CableLabs®
Access Network Independent

Standalone Router Specification

CL-SP-sRouter-I02-170111

ISSUED

DISCLAIMER

This document is furnished on an "AS IS" basis and neither CableLabs nor its members provides any representation or warranty, express or implied, regarding the accuracy, completeness, noninfringement, or fitness for a particular purpose of this document, or any document referenced herein. Any use or reliance on the information or opinion in this document is at the risk of the user, and CableLabs and its members shall not be liable for any damage or injury incurred by any person arising out of the completeness, accuracy, or utility of any information or opinion contained in the document.

CableLabs reserves the right to revise this document for any reason including, but not limited to, changes in laws, regulations, or standards promulgated by various entities, technology advances, or changes in equipment design, manufacturing techniques, or operating procedures described, or referred to, herein.

This document is not to be construed to suggest that any company modify or change any of its products or procedures, nor does this document represent a commitment by CableLabs or any of its members to purchase any product whether or not it meets the characteristics described in the document. Unless granted in a separate written agreement from CableLabs, nothing contained herein shall be construed to confer any license or right to any intellectual property. This document is not to be construed as an endorsement of any product or company or as the adoption or promulgation of any guidelines, standards, or recommendations.

Document Status Sheet

Document Control Number:	CL-SP-sRouter-I02-170111			
Document Title:	Standalone Router Specification			
Revision History:	I01 – Released 03/17/16 I02 – Released 01/11/17			
Date:	January 11, 2017			
Status:	~~Work in Progress~~	~~Draft~~	**Issued**	~~Closed~~
Distribution Restrictions:	~~Author Only~~	~~CL/Member~~	~~CL/ Member/ Vendor~~	**Public**

Key to Document Status Codes

Work in Progress	An incomplete document, designed to guide discussion and generate feedback that may include several alternative requirements for consideration.
Draft	A document in specification format considered largely complete, but lacking review by Members and vendors. Drafts are susceptible to substantial change during the review process.
Issued	A generally public document that has undergone Member and Technology Supplier review, cross-vendor interoperability, and is for Certification testing if applicable. Issued Specifications are subject to the Engineering Change Process.
Closed	A static document, reviewed, tested, validated, and closed to further engineering change requests to the specification through CableLabs.

Trademarks

Contents

Figures

Tables

1 SCOPE

This specification is part of the Access Network Independent family of specifications developed by Cable Television Laboratories, Inc. (CableLabs). This specification was developed for the benefit of the cable industry, and includes contributions by operators and vendors from North America, Europe, and other regions.

1.1 Introduction and Purpose

This specification defines a core set of features that enable multiple subscriber devices to gain access to operator provided high-speed data service independently of the underlying access network. This core set of features allows for both IPv4- and IPv6-enabled devices to gain connectivity to the Internet.

The core set of features defined in this specification includes the ability to provision multiple CPE devices, a description of how to forward data to and from CPE devices, and also the ability to forward IP Multicast traffic to and among CPE devices.

1.2 Requirements

Throughout this document, the words that are used to define the significance of particular requirements are capitalized. These words are:

"MUST"	This word means that the item is an absolute requirement of this specification.
"MUST NOT"	This phrase means that the item is an absolute prohibition of this specification.
"SHOULD"	This word means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before choosing a different course.
"SHOULD NOT"	This phrase means that there may exist valid reasons in particular circumstances when the listed behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.
"MAY"	This word means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

2 REFERENCES

2.1 Normative References

In order to claim compliance with this specification, it is necessary to conform to all or part of the following standards and other works as indicated, in addition to the other requirements of this specification. Notwithstanding, intellectual property rights may be required to use or implement such normative references.

[CANN DHCP-Reg]	CableLabs DHCP Options Registry Specification, CL-SP-CANN-DHCP-Reg-I14-170111, January 11, 2017, Cable Television Laboratories, Inc.
[CLAB-ANI-DEV-MIB]	CableLabs Access Network Independent Device, CLAB-ANI-DEV-MIB, http://www.cablelabs.com/MIBs/common/
[CLAB-GRE-MIB]	CableLabs Generic Route Encapsulation MIB, CLAB-GRE-MIB, http://www.cablelabs.com/MIBs/common/
[CLAB-MAP-MIB]	CableLabs Mapping of Address and Port MIB, CLAB-MAP-MIB, http://www.cablelabs.com/MIBs/common/
[CLAB-WIFI-MIB]	CableLabs Wireless MIB, CLAB-WIFI-MIB, http://www.cablelabs.com/MIBs/wireless/
[FIPS 140-2]	Federal Information Processing Standards Publication (FIPS PUB) 140-2, Security Requirements for Cryptographic Modules, June 2001.
[FIPS 180-4]	Federal Information Processing Standards Publication (FIPS PUB) 180-4, Secure Hash Standard, May 2014.
[ISO/IEC 29341]	ISO/IEC 29341-1-1:2011, Information technology -- UPnP Device Architecture -- Part 1-1: Universal Plug and Play Architecture Version 1.1, September 12, 2011.
[MULPIv3.0]	DOCSIS MAC and Upper Layer Protocol Interface Specification, CM-SP-MULPIv3.0-I30-170111, January 11, 2017, Cable Television Laboratories, Inc.
[OSSIv3.0]	DOCSIS Operations Support System Interface Specification, CM-SP-OSSIv3.0-I30-170111, January 11, 2017, Cable Television Laboratories, Inc.
[RFC 792]	IETF RFC 792, Internet Control Message Protocol, J. Postel, September 1981.
[RFC 826]	IETF RFC 826, An Ethernet Address Resolution Protocol, David C. Plummer, November, 1982.
[RFC 1122]	IETF RFC 1122, Requirements for Internet Hosts - Communication Layers, R. Braden, October, 1989.
[RFC 1812]	IETF RFC 1812, Requirements for IP Version 4 Routers, F. Baker, June 1995.
[RFC 1918]	IETF RFC 1918, Address Allocation for Private Internets, Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, E. Lear, February 1996.
[RFC 2131]	IETF RFC 2131, Dynamic Host Configuration Protocol, R. Droms, March, 1997.
[RFC 2132]	IETF RFC 2132, DHCP Options and BOOTP Vendor Extensions, S. Alexander, R. Droms, March 1997.
[RFC 2710]	IETF RFC 2710, Multicast Listener Discovery (MLD) for IPv6, S. Deering, W. Fenner, B. Haberman, October 1999.
[RFC 2784]	IETF RFC 2784, Generic Routing Encapsulation (GRE), D. Farinacci, T. Li, S. Hanks, D. Meyer, P. Traina, March 2000.
[RFC 2827]	IETF RFC 2827, Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing, P. Ferguson, D. Senie, May 2000.
[RFC 2863]	IETF RFC 2863, The Interfaces Group MIB, K. McCloghrie, F. Kastenholz, June 2000.
[RFC 2890]	IETF RFC 2890, Key and Sequence Number Extensions to GRE, G. Dommety, September 2000.

[RFC 3022] IETF RFC 3022, Traditional IP Network Address Translator (Traditional NAT), P. Srisuresh, K. Egevang, January 2001.

[RFC 3315] IETF RFC 3315, Dynamic Host Configuration Protocol for IPv6 (DHCPv6), R. Droms, Ed., J. Bound, B. Volz, T. Lemon, C. Perkins, M. Carney, July 2003.

[RFC 3376] IETF RFC 3376, Internet Group Management Protocol, Version 3, B. Cain, S. Deering, I. Kouvelas, B. Fenner, A. Thyagarajan, October, 2002.

[RFC 3418] IETF RFC 3418, Management Information Base (MIB) for the Simple Network Management Protocol (SNMP), R. Presuhn, December 2002.

[RFC 3633] IETF RFC 3633, IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6, O. Troan, R. Droms, December 2003.

[RFC 3646] IETF RFC 3646, DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6), R. Droms, December 2003.

[RFC 3736] IETF RFC 3736, Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6, R. Droms, April 2004.

[RFC 3810] IETF RFC 3810, Multicast Listener Discovery Version 2 (MLDv2) for IPv6, R. Vida, Ed., L. Costa, Ed., June 2004.

[RFC 4022] IETF RFC 4022, Management Information Base for the Transmission Control Protocol (TCP), R. Raghunarayan, March 2005.

[RFC 4113] IETF RFC 4113, Management Information Base for the User Datagram Protocol (UDP), B. Fenner, J. Flick, June 2005.

[RFC 4191] IETF RFC 4191, Default Router Preferences and More-Specific Routes, R. Draves, D. Thaler, November 2005.

[RFC 4242] IETF RFC 4242, Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6 (DHCPv6), S. Venaas, T. Chown, B. Volz, November 2005.

[RFC 4291] IETF RFC 4291, IP Version 6 Addressing Architecture, R. Hinden, S. Deering, February 2006.

[RFC 4292] IETF RFC 4292, IP Forwarding Table MIB, B. Haberman, April 2006.

[RFC 4293] IETF RFC 4293, Management Information Base for the Internet Protocol (IP), S. Routhier, (Editor), Bill Fenner, Brian Haberman, Dave Thaler, April 2006.

[RFC 4301] IETF RFC 4301, Security Architecture for the Internet Protocol, S. Kent, K. Seo, December 2005.

[RFC 4346] IETF RFC 4346, The Transport Layer Security (TLS) Protocol Version 1.1, T. Dierks, E. Rescorla, April 2006.

[RFC 4347] IETF RFC 4347, Datagram Transport Layer Security, E. Rescorla, N. Modadugu, April 2006.

[RFC 4361] IETF RFC 4361, Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4), T. Lemon, B. Sommerfeld, February 2006.

[RFC 4443] IETF RFC 4443, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification, A. Conta, S. Deering, M. Gupta, Ed., March 2006.

[RFC 4861] IETF RFC 4861, Neighbor Discovery for IP Version 6 (IPv6), T. Narten, E. Nordmark, W. Simpson, H. Soliman, September 2007.

[RFC 4862] IETF RFC 4862, IPv6 Stateless Address Autoconfiguration, S. Thomson, T. Narten, T. Jinmei, September 2007.

[RFC 4884] IETF RFC 4884, Extended ICMP to Support Multi-Part Messages, R. Bonica, D. Gan, D. Tappan, C. Pignataro, April 2007.

[RFC 5246] IETF RFC 5246, The Transport Layer Security (TLS) Protocol Version 1.2, T. Dierks, E. Rescorla, August 2008.

[RFC 5280] IETF RFC 5280, Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile, D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Hously, W. Polk, May 2008.

[RFC 5389] IETF RFC 5389, Session Traversal Utilities for NAT (STUN), J. Rosenberg. R, Mahy, P. Matthews, D. Wing, October 2008.

[RFC 5905] IETF RFC 5905, Network Time Protocol Version 4: Protocol and Algorithms Specification, D. Mills, J. Martin, J. Burbank, W. Kasch, June 2010,

[RFC 5908] IETF RFC 5908, Network Time Protocol (NTP) Server Options for DHCPv6, R. Gayroud, B. Lourdelet, June 2010,

[RFC 5942] IETF RFC 5942, IPv6 Subnet Model: The Relationship Between Links and Subnet Prefixes, H. Singh, W. Beebee, E. Nordmark, July 2010.

[RFC 6092] IETF RFC 6092, Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service, J. Woodyatt, Ed., January 2011.

[RFC 6106] IETF RFC 6106, IPv6 Router Advertisement Options for DNS Configuration, J. Jeong, S. Park, L. Beloeil, S. Madanapalli, November 2010.

[RFC 6333] IETF RFC 6333, Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion, A. Durand, R. Droms, J. Woodyatt, Y. Lee, August 2011.

[RFC 6334] IETF RFC 6334, Dynamic Host-Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite, D.Hankins, T. Mrugalski, August 2011.

[RFC 6347] IETF RFC 6347, Datagram Transport Layer Security Version 1.2, E. Rescorla, N. Modadugu, January 2012.

[RFC 6540] IETF RFC 6540, IPv6 Support Required for All IP-Capable Nodes. W. George, C. Donley, C. Liljenstolpe, L. Howard, April 2012.

[RFC 6762] IETF RFC 6762, Multicast DNS, S. Cheshire, M. Krochmal, Apple Inc., February 2013.

[RFC 7083] IETF RFC 7083, Modification to Default Values of SOL_MAX_RT and INF_MAX_RT, R. Droms, November 2013.

[RFC 7084] IETF RFC 7084, Basic Requirements for IPv6 Customer Edge Routers, H. Singh, W. Beebee, C. Donley, B. Stark, November 2013.

[RFC 7296] IETF RFC 7296, Internet Key Exchange Protocol Version 2 (IKEv2), C. Kaufman, P. Hoffman, Y. Nir, P. Eronen, T. Kivinen, October 2014.

[RFC 7599] IETF RFC 7599, Mapping of Address and Port Using Translation (MAP-T), X. Li, C. Bao, W. Dec, O. Troan, S. Matsushima, T. Murakami, July 2015.

[RSA 1] RSA Laboratories, PKCS #1: RSA Encryption Standard. Version 1.5, RSA Security, Inc., Bedford, MA, November 1993.

[SECv3.0] DOCSIS Security Specification, CM-SP-SECv3.0-I16-160602, June 2, 2016, Cable Television Laboratories, Inc.

[TR-069] TR-069 CPE WAN Management Protocol v1, Issue 1 Amendment 5, November 2013, Broadband Forum Technical Report.

[TR-181] TR-181 Device Data Model for TR-069, Issue 2 Amendment 8, September 2014, Broadband Forum Technical Report.

2.2 Informative References

This specification uses the following informative references.

[MC-EMC] IP Multicast Controller-Client Interface Specification, OC-SP-MC-EMCI-C01-161026, October 26, 2016, Cable Television Laboratories, Inc.

[MR-230]	TR-069 Deployment Scenarios, Issue 1, MR-230, August 2010, Broadband Forum Marketing Report.
[RFC 793]	IETF RFC 793/STD-7, Transmission Control Protocol, J. Postel, September 1981.
[RFC 1323]	IETF RFC 1323, TCP Extensions for High Performance, V. Jacobson, B. Braden, and D. Borman, May 1992.
[RFC 2460]	IETF RFC 2460, Internet Protocol, Version 6 (IPv6) Specification, S. Deering, and R. Hinden, December 1998.
[RFC 3775]	IETF RFC 3775, Mobility Support in IPv6, D. Johnson, C. Perkins, and J. Arkko, June 2004.
[RFC 3828]	IETF RFC 3828, The Lightweight User Datagram Protocol (UDP-Lite), L-A. Larzon, M. Degermark, S. Pink, L-E. Jonsson, and G. Fairhurst, July 2004.
[RFC 3879]	IETF RFC 3879, Deprecating Site Local Addresses, C. Huitema, B. Carpenter, September 2004.
[RFC 4007]	IETF RFC 4007, IPv6 Scoped Address Architecture, S. Deering, B. Haberman, T. Jinmei, E. Nordmark, and B. Zill, March 2005.
[RFC 4193]	IETF RFC 4193 Unique Local IPv6 Unicast Addresses. R. Hinden, B. Haberman, October 2005.
[RFC 4302]	IETF RFC 4302, IP Authentication Header, S. Kent, December 2005.
[RFC 4303]	IETF RFC 4303, IP Encapsulating Security Payload (ESP). S. Kent, December 2005.
[RFC 4340]	IETF RFC 4340, Datagram Congestion Control Protocol (DCCP), E. Kohler, M. Handley, and S. Floyd, March 2006.
[RFC 4960]	IETF RFC 4960, Stream Control Transmission Protocol, R. Stewart, September 2007.
[RFC 5095]	IETF RFC 5095, Deprecation of Type 0 Routing Headers in IPv6, J. Abley, P. Savola, and G. Neville-Neil, December 2007.
[RFC 5156]	IETF RFC 5156, Special-Use IPv6 Addresses, M. Blanchet, April 2008.
[RFC 5201]	IETF RFC 5201, Host Identity Protocol, R. Moskowitz, P. Nikander, P. Jokela, and T. Henderson, April 2008.
[RFC 5382]	IETF RFC 5382, NAT Behavioral Requirements for TCP, S. Guha, K. Biswas, B. Ford, S. Sivakumar, and P. Srisuresh, October 2008.
[RFC 5996]	IETF RFC 5996, Internet Key Exchange Protocol Version 2 (IKEv2), C. Kaufman, P. Hoffman, Y. Nir, and P. Eronen, September 2010.
[TR-106]	TR-106 Data Model Template for TR-069-Enabled Devices, Issue 1, Amendment 7, September 2013, Broadband Forum Technical Report.
[WI-FI MGMT]	Wi-Fi Provisioning Framework Specification, WR-SP-WiFi-MGMT-I07-161213, December 13, 2016, Cable Television Laboratories, Inc.

2.3 Reference Acquisition

- Broadband Forum, 48377 Fremont Blvd, Suite 117 Fremont, CA 94538 USA;
  Phone: +1-510-492-4020; Fax: +1-510-492-4001; http://www.broadband-forum.org
- Cable Television Laboratories, Inc., 858 Coal Creek Circle, Louisville, CO 80027 USA;
  Phone:+1-303-661-9100; Fax:+1-303-661-9199; http://www.cablelabs.com
- Internet Engineering Task Force (IETF) Secretariat, 48377 Fremont Blvd., Suite 117, Fremont, CA 94538 USA;
  Phone: +1-510-492-4080; Fax: +1-510-492-4001, http://www.ietf.org/
- International Organization for Standardization, http://www.iso.org/iso

3 TERMS AND DEFINITIONS

This specification uses the following terms:

Access Network Demarcation Device	The access network endpoint CPE (e.g., DOCSIS CM or EPON ONU) that provides the interface(s) to the managed equipment in the customer's site.
Best Current Practice	A term typically associated with recommendations and numbered standards provided by the Internet Engineer Task Force (IETF)
Customer Edge Router	Provides specific services and forwarding capabilities necessary for establishing and maintaining the customer edge on the Operator-Facing Interface (WAN). In this role, sRouter application services such as DHCP, NAPT and Packet Filtering Firewall are enabled.
Customer-Facing Interface	An sRouter interface used for connecting CPE devices. Defined in [RFC 7084] as a Local Area Network (LAN) Interface, the Customer-Facing Interface is represented by a physical port.
Customer-Facing IP Interface	An IP Interface is a connected interface to the sRouter that is not necessarily mapped one-to-one with the number of Customer-Facing ports on the sRouter. As defined in [RFC 7084], this is an IP LAN interface in which one or many physical ports are associated with an IP address.
Customer-Facing Logical Interface	A logical Interface connected to the sRouter that is not necessarily mapped one-to-one with the number of Customer-Facing ports on the sRouter. As defined in [RFC 7084], this is a LAN interface in which one or more physical ports are associated with a logical interface, such as a VLAN.
Down Interface	An interface on an sRouter that is topologically further away from the Operator's network than the 'Up' interface on that same router. See Up Interface.
Hard Reset	Describes a full reset of the device.
Link ID	16 bits of both IPv4 and IPv6 addresses chosen to uniquely identify each "link" or LAN segment (Customer-Facing IP Interface) within the home network. Counting from the left, the Link ID includes bits 49 - 64 (fourth 16-bit block) in an IPv6 address and bits 9 - 24 (middle two octets) in an IPv4 address.
Multicast Subscription Database	A simple table of entries for the IPv4 or IPv6 Multicast Group Membership information maintained by the sRouter on respective interfaces. Implementation details for storage of records are completely vendor-defined.
Operator-Facing Interface	The IP Router interface that is connected to the access network demarcation device such as a CM or ONU. As defined in [RFC 7084], this is a Wide Area Network (WAN) interface. In CPE WAN Management Protocol (CWMP) this is called an upstream interface.
Operator-Facing IP Interface	IP Interface that is provisioned with an IP Address provided by the Operator. As defined in [RFC 7084], this is a WAN interface.
Prefix	A common address component, which defines a portion of a network. The meanings of the terms "prefix" and "subnet" are used interchangeably. The term prefix is favored in this document. See also PD.
Prefix Delegation	A form of IPv6 address assignment allowing the operator's DHCP server to *delegate* a prefix of a specific length, such as /56, to a customer's Router. The delegation of one or more prefixes allows the Router to further sub-divide and assign individual prefixes (which are /64 in length) to its interfaces and/or provide prefix sub-delegation to additional Routers within the customer's network. Prefix Delegation occurs only between the operator's DHCP server and a Router operating in the role of Customer Edge Router (CER). See Prefix Sub-Delegation. See also Customer Edge Router.

Reset	Describes a routine in which the operational state is interrupted by the instruction to shut down and restart. The term is synonymous with the terms re-initialization and reboot.
Router Advertisement	A periodic message containing one or more of the following: Prefix information including prefix lifetime, default router and other information.
Router Information Option	An encoding within the Router Advertisement containing next hop and Default Router information
Router Solicitation	A client message requesting the Router Advertisement
Service Discovery	A set of protocols and methods that is used to discover services that are made available by hosts and nodes within the customer network.
sRouter	A standalone router that conforms to the requirements in this specification. It is intended to be network access independent with regard to provisioning and management.
Subnet	A portion of a network that shares a common address component. The meanings of the terms "prefix" and "subnet" are used interchangeably. The term "prefix" is favored in this document.
TR-069	Term used to refer to the CPE WAN Management Protocol suite defined in [TR-069].
TR-069 CPE	Term used to refer to the CPE managed using the CPE WAN Management Protocol suite defined in [TR-069].
Up Interface	A router interface that connects to another router that is closer to the ISP network. For example, the ‘Up Interface’ of an Internal router is the port used to connect to the CFI (Down interface), of the Router. See Down Interface.

4 ABBREVIATIONS AND ACRONYMS

This specification uses the following abbreviations:

ACS	Auto-configuration Server
AFTR	Address Family Translation Router
AH	Authentication Header
ALG	Application Layer Gateway
ARP	Address Resolution Protocol
B4	Basic-Bridging Broadband
BCP	Best Current Practice
BMR	Basic Mapping Rule
BR	Border Relay
CER	Customer Edge Router
CER-ID	Customer Edge Router-Identification
CM	Cable Modem
CMTS	Cable Modem Termination System
CPE	Customer Premises Equipment
CWMP	CPE WAN Management Protocol
DAD	Duplicate Address Detection
DCCP	Datagram Congestion Control Protocol
DHCPv4	Dynamic Host Configuration Protocol version 4
DHCPv6	Dynamic Host Configuration Protocol version 6
DNS	Domain Name Service
DUID	DHCP unique identifier
DUID-EN	DUID Enterprise Number
DUID-LL	DUID Link Layer address
DUID-LLT	DUID Link Layer plus Time
EA	Embedded Address
ESP	Encapsulating Security Protocol
EUI	Extended Unique Identifier
FMR	Forwarding Mapping Rule
FTP	File Transfer Protocol
GNAP	Global Network Address Port
GRE	Generic Route Encapsulation
GUA	Global Unique Address
GW	Gateway
IA_NA	Identity Association for Non-temporary Addresses
IA_PD	Identity Association for Prefix Delegation
IANA	Internet Assigned Numbers Authority
ICMP	Internet Control Message Protocol
ID	Identifier

IGMP	Internet Group Management Protocol
IKE	Internet Key Exchange
IPTV	Internet Protocol Television
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
IRT	Initial Retransmission Times
LAN	Local Area Network
MAC	Media Access Control
MAP-E	Mapping of Address and Port (Encapsulation)
MAP-T	Mapping of Address and Port (Translation)
mDNS	Multicast Domain Name System
MIB	Management Information Base
MLD	Multicast Listener Discovery
MoCA	Multimedia over Coax Alliance
MRC	Maximum Retransmission Count
MRD	Maximum Retransmission Duration
MRT	Maximum Retransmission Time
MSS	Maximum Segment Size
MTU	Maximum Transmission Unit
NAPT	Network Address Port Translation
NAT	Network Address Translation
ND	Neighbor Discovery
NS	Neighbor Solicitation
NTPv4	Network Time Protocol version 4
OID	Object ID
ONU	Optical Network Unit
ORCHIDv2	Overlay Routable Cryptographic Task Identifiers version 2
ORO	Option Request Option (DHCP)
OUI	Organization Unique Identifier
PD	Prefix Delegation
PDU	Protocol Data Unit
PIO	Prefix Information Option
PNAP	Private Network Address Port
QoS	Quality of Service
RA	Router Advertisement
RD	Router Discovery
RFC	Request For Comment
RG	Residential Gateway
RIO	Router Information Option.
RS	Router Solicitation
SCTP	Stream Control Transmission Protocol

SIP	Session Initiation Protocol
SLAAC	Stateless Address Autoconfiguration
SNMP	Simple Network Management Protocol
sRouter	Standalone Router
SSDP	Simple Service Delivery Protocol
TCP	Transmission Control Protocol
ToS	Type of Service
TTL	Time To Live
UDP	User Datagram Protocol
ULA	Unique Local Addresses
USB	Universal Serial Bus
VLAN	Virtual Local Area Network
WAN	Wide Area Network

5 THEORY OF OPERATION

This specification defines the following:

A) Provisioning of the IP-enabled modes of the sRouter
B) The sRouter serving as a Customer Edge Router
C) The provisioning of the CPE devices with IPv4 and IPv6 addressing
D) IPv6 data forwarding, as well as IPv4 data forwarding with and without NAPT
E) IP multicast traffic forwarding
F) The preservation of IP QoS markings applied to IP data forwarded to and from the CPE devices
G) GRE tunnel configuration and management
H) Service Discovery

This specification defines requirements for an sRouter device with a single Operator-Facing IP Interface and one or more Customer-Facing IP Interfaces. The sRouter makes use of [TR-069] CWMP for its Operator-Facing management interface options for the provisioning and management of the sRouter. In addition, this specification defines configuration and monitoring via SNMP. However, SNMP configuration is distinguished from TR-069 in that SNMP SET PDUs can only be sent *after* the sRouter has become operational.

This specification defines two methods that the sRouter uses to assign IP addresses to its Customer-Facing Interfaces. These two methods include Link-ID and non-Link-ID addressing schemes. Both methods implement IPv4 addressing using [RFC 1918] IPv4 address ranges, but only the Link-ID method uses an algorithm to derive a unique IPv4 subnet to support multiple routers within the customer LAN. Link-ID provides a predictable IPv4 addressing scheme, where the IPv6 link bits obtained during IPv6 prefix acquisition are reflected in IPv4 octets 2 and 3, and enabling native IPv4 routing within the customer network. This functionality allows for routing across several routers without routing protocols or the encumbrance of two or more layers of NAPT. However, if Link-ID is not enabled or an IPv6 prefix is not available from which to generate a Link-ID, then IPv4 routing across multiple routers will not be possible unless the sRouter is manually configured.

The primary functions provided by the sRouter defined by this specification include connecting multiple customer CPE devices to the operator's high-speed Internet service, and provides essential network application services to support the customer's LAN. The sRouter is delegated the responsibility of provisioning CPE devices in the customer's LAN, acting in the role of the Customer Edge Router (CER). Typical network services provided by the sRouter include DHCP, packet filtering firewall, and NAPT.

In addition, the sRouter may be configured to support IPv4 only, IPv6 only and simultaneous IPv4/IPV6 Operator-Facing IP Interfaces as defined in Table 1. By default, the sRouter assumes Dual Protocol Enabled mode operation in the absence of explicit provisioning information to the contrary.

Table 1 - IP-enabled Modes

IP Enabled Mode	Forwarding Behavior
IPv4 Protocol Enabled mode	IPv4 Route Forwarding
IPv6 Protocol Enabled mode	IPv6 Route Forwarding
Dual Protocol Enabled mode	Both IPv4 and IPv6 Route Forwarding

5.1 Operational Use Case

The following sub-section describes the use cases for the sRouter when operating in a layer 3 forwarding mode. In each of the operational use cases, the following assumptions also apply:

- Community Wi-Fi Support, provisioning and management of the public and private SSIDs

5.1.1 sRouter as the Customer Edge Router

This use case is characterized by a single sRouter providing services and traffic forwarding within the customer's network core over a single, operator managed access network. A single sRouter operates as the Customer Edge Router, which provides GRE tunnel origination to the operator's Wi-Fi core for Community Wi-Fi services, Internet

services, DHCP, packet filtering firewall and, depending on the specific IP Protocol Enabled mode, NAPT services. It is assumed that a single IA_PD is assigned by the operator with sufficient length to allow the router to assign prefixes to interfaces

NOTE: Community Wi-Fi SSIDs are under the control of the operator's Wi-Fi core, not the sRouter.

Table 2 - sRouter Network Roles

sRouter Role in the Network	Associated Functions and Services
Customer Edge Router (CER)	IPv4 Services: • Layer 3 forwarding • NAPT • DHCPv4 • Packet Filtering Firewall • DNS Forwarder IPv6 Services: • Layer 3 forwarding • DHCPv6 • DHCP-PD • Packet Filtering Firewall [RFC 6092] • DNS Forwarder Community Wi-Fi: • GRE tunnel origination and public/private SSID differentiation • Wi-Fi MGMT functions, such as RADIUS AAA

5.2 TR-069 Architecture

This section defines TR-069 requirements for the sRouter management architecture.

The TR-069 specification suite defines the Device 2.x data model in [TR-181]. It refers to a CPE device management space for holding the device itself and root of other service specifications (e.g., VoIP, storage, and IPTV). See [MR-230] for more details on TR-069 deployment scenarios. TR-069 is access network technology agnostic.

TR-069 allows the transparent integration and management of access network technologies within the sRouter and CPE Services by combining the components and their respective management data. Figure 1 is based on the "Simple Router Example (Interfaces Visualized)" figure of [TR-181]. In Figure 1, the stack layers are seen as interfaces per [TR-181], physical interfaces (e.g., Ethernet, SSID, Wi-Fi Radio), bridges, and IP Interfaces.

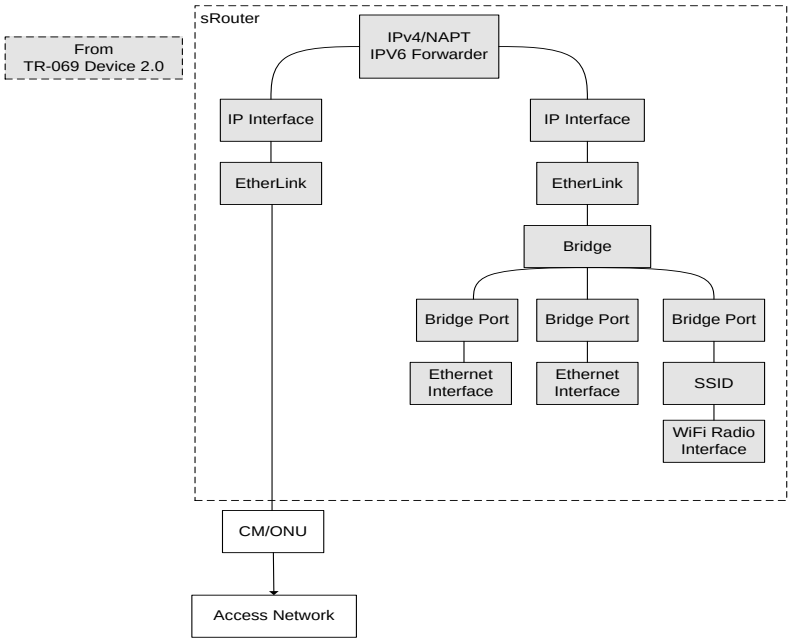

Figure 1 - TR-069 Interface Model Applied to sRouter

5.3 sRouter Device Management

The sRouter enables remote management by supporting TR-069 and optionally, SNMP.

5.4 Service Discovery

Service Discovery will allow LAN-attached devices with services to announce their presence and allow a query/response method for discovering and choosing a service from a list of possible candidates that provide that service. For example, a network-based printer could be discovered by one or more service discovery techniques. The mDNS protocol, as defined in [RFC 6762], and UPnP, as defined by the UPnP Forum and specified in [ISO/IEC 29341], is used for providing Service Discovery in the sRouter.

5.4.1 mDNS (multicast Domain Name System)

The mDNS protocol provides both an announcement and a query/response mechanism to provide a list of devices that offer services on the home network. The mDNS protocol is link-scoped but can be enhanced to provide service to multiple networks in the customer network. A method to relay announcements and queries/responses between different networks is therefore needed.

The sRouter takes an announcement, query or response packet from one link/subnet and relays it onto another link/subnet, but replaces the IP source address from the originating link/subnet with the sRouter's egress IP address on the other link/subnet. Additionally, if the payload of the mDNS packet contains a link-local or Auto-IP resource record, those address records are removed before the packet is placed onto the new link/subnet.

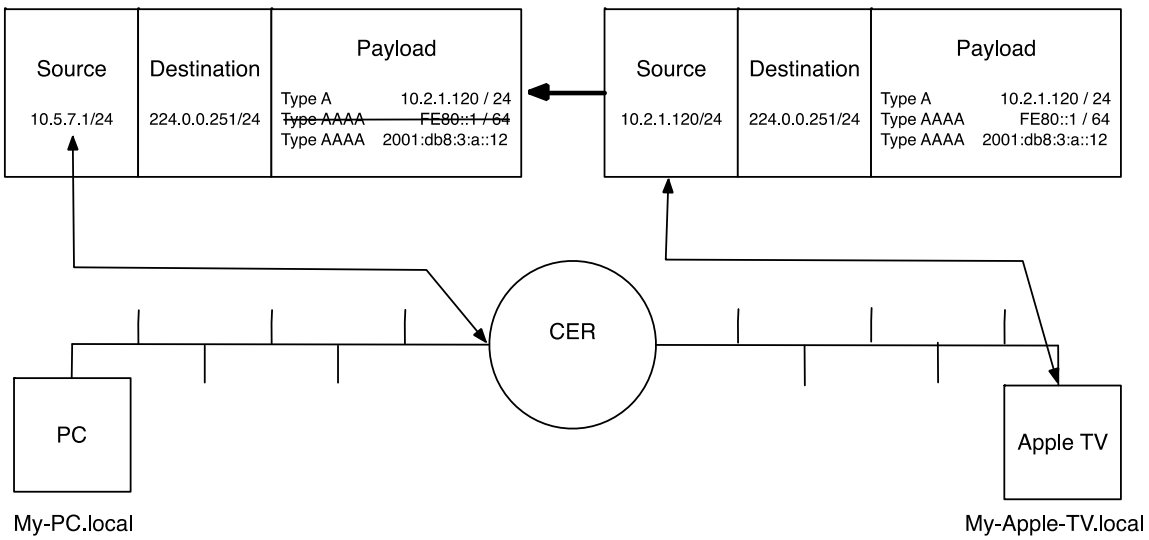

Figure 2 - mDNS Source IP and Payload Changes Going Through an sRouter

Hosts implementing the mDNS protocol might silently discard any frame with a source IP address that is not part of the receiver's network. When the sRouter receives a packet with an address that is not on the receiver's network, the sRouter MUST replace the source address in the packet with the IP address of the sRouter's egress interface to the receiver's network. This ensures that the packet is properly relayed to other hosts and not dropped by the host.

mDNS forwards as described in Figure 2 above to multicast addresses (FF02::FB and 224.0.0.251).

The sRouter will not keep state information, but simply relays the packet and makes the appropriate changes to the IP source address and payload. The sRouter MUST NOT forward or listen for announcement, query or response messages, for either IPv4 or IPv6, on the Operator-Facing Interface.

5.4.2 UPnP (Universal Plug and Play)

The UPnP architecture makes use of a number of protocols including IP, TCP, UDP, HTML, and XML to enable peer-to-peer networking and service discovery. Most of the protocols used by UPnP will function across network segments in the home network. The one exception and the focus of this section is the UPnP Device Discovery Protocol, which uses the Simple Service Discovery Protocol (SSDP). UPnP was designed to function in an unmanaged network environment by having UPnP controllers (control points) automatically discover UPnP devices. SSDP is used by UPnP controllers to search for UPnP devices and is also used by UPnP devices to announce themselves and their services to UPnP controllers.

UPnP forwards to multicast addresses (FF02::C, FF05::C, FF02::130 and 239.255.255.250) as shown in Figure 3. The sRouter is expected to forward to the site scoped address only.

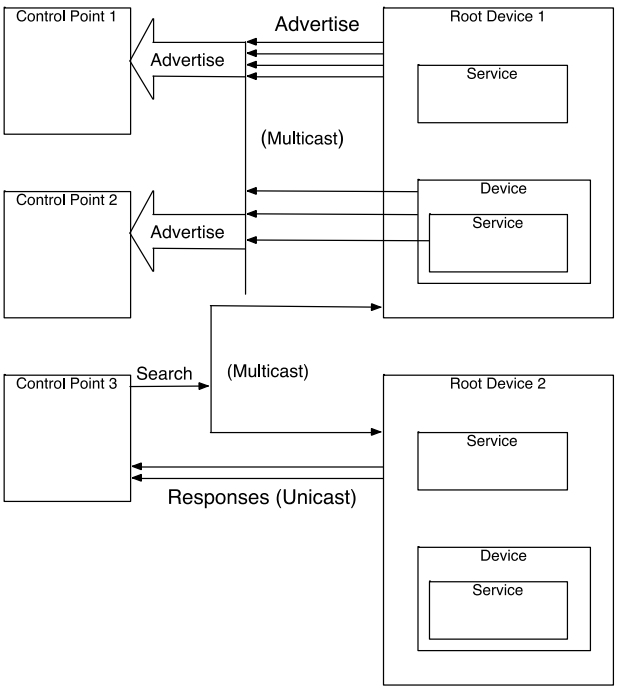

Figure 3 - UPnP General Architecture

The sRouter does not forward or listen for SSDP messages, over either IPv4 or IPv6, on the Operator-Facing Interface. IPv6 support was added to annex A of the [ISO/IEC 29341].

5.5 CER-ID (Customer Edge Router)

An sRouter is situated at the edge of the customer network facing the operator's network, and acts as the demarcation point. When the sRouter is operating in this role, it is referred to as the CER. The CER-ID is used by routers to determine their role in the LAN, which may impact the services the router should enable. The CER-ID is passed from DHCP server to connected sRouters which then compare the information in the CER-ID and the sRouter's own WAN interface addressing to determine its role. The sRouter compares the IA_NA address assigned to its Customer-Facing IP Interface to the assigned IA_PD on the Customer-Facing IP Interface. If the two address scopes are derived from the same parent prefix, then the router determines that it is not the CER, but is a router operating within the interior of the LAN. For the purposes of this specification, the only role currently supported is the CER and is assumed regardless of the settings of the CER-ID or the absence of the CER-ID from the DHCP response.

5.6 IPv4 Multicast NAPT

IPv4 Multicast packets destined to Multicast servers outside the customer's home network are a special case for NAPT and need special handling at the sRouter. One scenario where forwarding of IP Multicast packets at the sRouter needs special handling is when a video source is using a private network address on a Customer-Facing IP Interface. In general, for video sources on the Customer-Facing IP Interface to work, the sRouter is required to run at least one industry-standard multicast routing protocol to advertise the flows.

Since the sRouter supports IGMP proxy for IGMP v2 and v3, there is no reason to support a special translation for multicast packets in the sRouter for IGMP messages from private network addresses arriving on the Customer-Facing IP Interface, as they are consumed by the sRouter and new IGMP messages are sent by the proxy agent from a public source network address on the Operator-Facing IP Interface.

As there is no NAPT for IPv6 Multicast traffic, the sRouter will forward MLD IPv6 Multicast traffic that is not link or site scoped through the Operator Facing Interface that originates from LAN based devices.

6 sRouter INITIALIZATION

The sRouter operates in any one of three possible modes – 'IPv4 Protocol Enabled', 'IPv6 Protocol Enabled', or 'Dual IP Protocol Enabled', as summarized in Table 3, and the sRouter MUST support all three modes of operation. The sRouter MUST default to 'Dual IP Protocol Enabled' mode in conformance with [RFC 6540].

The sRouter initializes, by default, in 'Dual IP Protocol Enabled' mode. After initialization, the operator may change the mode to any of possible modes (e.g., 'IPv4 Protocol Enabled', 'IPv6 Protocol Enabled') using either TR-069 or SNMP.

The sRouter IP Protocol Mode MUST persist across initializations.

When the sRouter is in 'IPv4 Protocol Enabled' mode, the sRouter performs IPv4 provisioning as described in Section 7 and IPv4 data forwarding and NAPT according to Section 9. The sRouter operating in 'IPv4 Protocol Enabled' mode does not perform any IPv6 provisioning. When the sRouter is in 'IPv4 Protocol Enabled' mode, the sRouter MUST NOT forward IPv6 traffic between the Operator-Facing Interface and the Customer-Facing Interfaces.

When the sRouter is in 'IPv6 Protocol Enabled' mode, the sRouter performs IPv6 provisioning according to Section 8 and IPv6 data forwarding according to Section 10. The sRouter operating in 'IPv6 Protocol Enabled' mode does not perform any IPv4 provisioning. When the sRouter is in IPv6 Protocol Enabled' mode, the sRouter MUST NOT forward IPv4 traffic between the Operator-Facing Interface and the Customer-Facing Interfaces.

When the sRouter is in 'Dual IP Protocol Enabled' mode, the sRouter performs IPv4 provisioning as described in Section 7 and IPv6 provisioning according to Section 8. Once an sRouter in 'Dual IP Protocol Enabled' mode acquires an IPv6 address and prefix per Section 8, the sRouter performs IPv6 data forwarding according to Section 10. An sRouter in 'Dual IP Protocol Enabled' mode acquires an IPv4 address per Section 7 and then performs IPv4 data forwarding and NAPT according to Section 9.

When the sRouter is enabled in any of the IP Protocol Enabled Modes, the sRouter MUST forward IP traffic between the Customer-Facing Interfaces, regardless of which IP Protocol Mode is enabled.

Table 3 provides a summary of the sRouter behavior based upon the configured mode of operation as well as when it is 'Disabled'.

Table 3 - Operator-Facing sRouter Modes

Mode	IPv4	IPv6
IPv4 Protocol Enabled	IPv4 Provisioning (Section 7). IPv4 data forwarding using NAPT (Section 9).	No IPv6 provisioning. No IPv6 data forwarding between Operator-Facing Interface and the Customer-Facing Interfaces.
IPv6 Protocol Enabled	No IPv4 provisioning. No IPv4 data forwarding between Operator-Facing Interface and the Customer-Facing Interfaces.	IPv6 Provisioning (Section 8). IPv6 data forwarding (Section 10).
Dual IP Protocol Enabled	IPv4 Provisioning (Section 7). IPv4 data forwarding using NAPT (Section 9).	IPv6 Provisioning (Section 8). IPv6 data forwarding (Section 10).

6.1 Network Time Protocol

Network Time Protocol version 4 (NTPv4) is used to provide time synchronization to the sRouter from one or more master servers. Such time synchronization is essential for correlating events in the sRouter's local log, and for providing accurate time for features that, while not explicitly defined in the sRouter specification, are commonplace in current implementations. These features and applications include content/parental controls, video content controls for such features as IP video and Digital Video Recorder, voice applications such as caller ID, and voicemail notification and similar applications that rely on a standard time reference to perform specific actions or provide a specific date and time in the application's interface.

The sRouter implements the NTPv4 protocol client per [RFC 5905] with the following exceptions:

- The sRouter MUST act as a client per section 2, Modes of Operation using the Operator Facing Interface.
- The sRouter MUST support a minimum of three NTP servers in the NTP server list.
- The sRouter MAY support the broadcast protocol mode of NTPv4 per section 3, Protocol Modes, with the provision that a filtering mechanism is provided to prevent security vulnerabilities as described in section 15, Security Considerations.
- The sRouter MAY support dynamic server discovery via multicast or anycast mechanisms per section 3.1, Dynamic Server Discovery.
- The sRouter MUST support the client/server protocol mode per section 3, Protocol Modes.

The sRouter implements the NTPv4 protocol Server for LAN Clients per [RFC 5905] with the following exceptions:

- The sRouter MUST populate the DHCP scopes defined for each Customer-Facing Interface with the NTPv4 server information obtained from the Operator Facing Interface IP provisioning.

7 IPV4 PROVISIONING

The normative requirements of this section are mandatory for an sRouter operating in 'IPv4 Protocol Enabled' mode and/or the 'Dual IP Protocol Enabled' mode as defined in Section 6.

The sRouter MUST use DHCPv4 [RFC 2131] via its Operator-Facing Interface in order to obtain an IP address and any other parameters needed to establish IP connectivity, as illustrated in Figure 4.

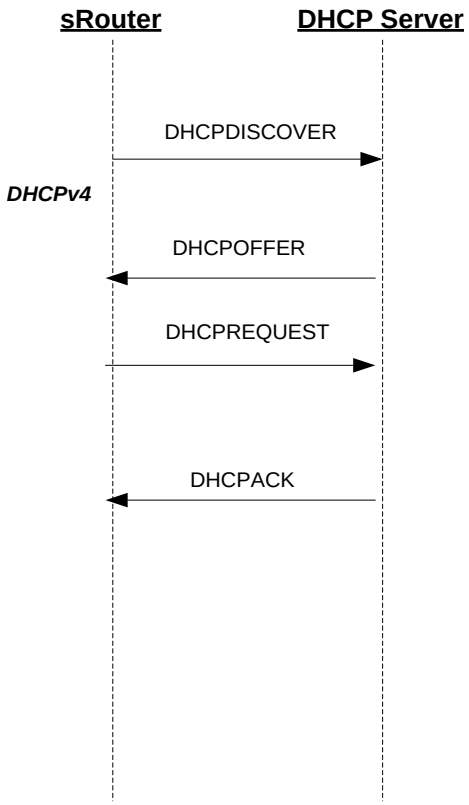

Figure 4 - IPv4 Provisioning Message Flow

The sRouter can receive multiple DHCPOFFER messages in response to its DHCPDISCOVER message. If a received DHCPOFFER message does not include all of the required DHCPv4 fields and options as described in Section 7.2, the sRouter MUST discard the DHCPOFFER message and wait for another DHCPOFFER message. If none of the received DHCPOFFER messages contain all the required DHCPv4 fields and options, the sRouter MUST retransmit the DHCPDISCOVER message.

The backoff values for the sRouter retransmission of DHCPDISCOVER messages SHOULD be chosen according to a uniform distribution between the minimum and maximum values in the rows of Table 4.

Table 4 - sRouter DHCP Retransmission Interval

Backoff Number	Minimum (sec.)	Maximum (sec.)
1	3	5
2	7	9
3	15	17
4	31	33
5	63	65

The sRouter SHOULD also implement a different retransmission strategy for the RENEWING and REBINDING states, as recommended in [RFC 2131], which is based on one-half of the remaining lease time.

The sRouter MUST limit the number of retransmissions of the DHCPDISCOVER and DHCPREQUEST messages to five or fewer. The sRouter MUST NOT forward IPv4 traffic between its Customer-Facing Interface and its Operator-Facing Interface until it has completed IPv4 provisioning, including the successful receipt of a DHCPACK message. The sRouter MUST NOT forward IPv4 traffic if, at any time, it does not have an IPv4 address for its Operator-Facing Interface.

The sRouter MUST be able to accept a unicast response from the DHCP server/relay agent.

7.1 DHCPv4 Fields Used by the sRouter

The sRouter MUST include the following fields in the DHCPDISCOVER and DHCPREQUEST messages:

- The hardware type (htype) MUST be set to 1.
- The hardware length (hlen) MUST be set to 6.
- The client hardware address (chaddr) MUST be set to the 48-bit MAC address associated with the IPv4 Operator-facing interface of the sRouter.
- The broadcast bit MUST NOT be set.
- The client-identifier option MUST be included, using the format defined in [RFC 4361].
- The parameter request list option MUST be included.
- The following option codes (defined in [RFC 2132] and [RFC 4361]) MUST be included in the list:
  - Option code 1 (Subnet Mask)
  - Option code 3 (Router Option)
  - Option code 6 (DNS Server Option)
  - Option code 42 Network Time Protocol Servers Option
  - Option code 55 (Parameter Request List)
  - Option code 60 (Vendor Class Identifier) [sRouter 1.0]

The following fields are expected in the DHCPOFFER and DHCPACK messages returned to the sRouter. The sRouter MUST configure itself with the listed fields from the DHCPACK:

- The IP address to be used by the sRouter (yiaddr) (critical).
- The IP Address lease time, option 51 (critical).
- The Server identifier, option 54 (critical).
- The subnet mask to be used by the sRouter (Subnet Mask, option 1) (critical).
- A list of addresses of one or more routers is to be used for forwarding sRouter-originated IP traffic (Router Option, option 3 (critical).
- A list of DNS Server addresses (critical).

NOTE: The sRouter is not required to use more than one router IP address for forwarding.

If a critical field is missing or invalid in the DHCPACK received during initialization, the sRouter MUST restart the DHCP cycle, beginning with an initial DHCPDISCOVER.

If a non-critical field is missing or invalid in the DHCPACK received during initialization, the sRouter MUST ignore the field, and continue the provisioning process.

If the yiaddr, Server Address, or Lease Time field is missing or invalid in the DHCPACK received during a renew or rebind operation, the sRouter MUST retry the renew or rebind operation until either: (1) it receives a response containing valid values of the yiaddr, Server Address, and Lease Time fields; or (2) the lease expires. If the lease expires, the sRouter MUST restart the DHCP cycle, beginning with an initial DHCPDISCOVER.

If any field other than the yiaddr, Server Address or Lease Time is missing, or is invalid in the DHCPACK received during a renew or rebind operation, the sRouter MUST ignore the field if it is invalid and remain operational.

7.2 sRouter DHCPv4 Server Sub-element

The DHCP server is responsible for assigning network address leases to LAN IP devices associated with Customer-Facing Interfaces. It is also responsible for providing LAN IP devices with configuration information via DHCP Option codes, as specified in [RFC 2132].

7.2.1 DHCPv4 Server Function Goals

Goals for the DHCP server include the following:

- Assign network address leases to CPE devices according to [RFC 2131].
- Assign private CPE addresses according to [RFC 1918].
- Assign configuration information according to [RFC 2132].

7.2.2 DHCPv4 Server Function System Description

The sRouter DHCPv4 server responsibilities include:

- Assigning IP Addresses and delivering DHCP configuration parameters to CPE Devices. The server relies on built-in default values for initial IP Address pool configuration, lease parameter configuration, and DHCP options values.
- Optional logging of DHCPv4 server errors to a local event log.

7.2.3 DHCPv4 Server Function Requirements

The sRouter MUST include a DHCPv4 server compliant with [RFC 2131].

In addition, the following requirements apply to the DHCPv4 Server function:

- When the sRouter DHCP server assigns an active lease for an IP address to a CPE Device, the server function of the sRouter MUST remove that IP address from the pool of IP addresses available for assignment.
- The requirements in this section use an appropriate address space as defined in [RFC 1918], and overrides the requirements in Section 8.9 when using IPv4-only mode (not Dual-Stack).
- The DHCP server function of the sRouter MUST support the DHCP options indicated as mandatory in Table 5.
- The DHCP server function of the sRouter MUST NOT respond to DHCP messages that are received through the Operator-Facing Interface, nor originate DHCP messages from the Operator-Facing Interface.
- The DHCP server function of the sRouter MUST NOT deliver any DHCP option with null value to any CPE device.
- The DHCP server function of the sRouter SHOULD be operational independent of the sRouter Operator-Facing Interface connectivity state.
- If the sRouter Operator-Facing Interface is not successfully provisioned, the sRouter DHCP server function SHOULD assign a short lease time to CPE devices and may omit options it has not acquired.
- The DHCP server function of the sRouter MUST assign private IP address space as defined in [RFC 1918].
- The DHCP server function of the sRouter SHOULD log errors to a local event log.

Table 5 - DHCPv4 Server Options

Option Number	Option Function
0	Pad
255	End
1	Subnet Mask
3	Router Option
6	Domain Name Server
42	Network Time Protocol Servers Option
50	Requested IP Address
51	IP Address Lease Time
54	Server Identifier
55	Parameter Request List
4491.3	Option(s) acquired under CL_V4EROUTER_CONTAINER_OPTION from the Operator

7.3 Operator-Facing IPv4 Address Release Behavior

There are a number of situations in which it is desirable for the sRouter to release its associated IPv4 address leases in order to protect the integrity of the DHCP database. Examples of such circumstances include situations in which the sRouter needs to be administratively reset (i.e., for configuration change, software update, or other reasons), or a change to the IPv4 address during DHCPv4 renewal.

Whenever the sRouter is instructed to reset, the sRouter MUST send a DHCP_RELEASE message [RFC 2131] for the IPv4 public address assigned by the DHCPv4 server to the sRouter's Operator-Facing Interface. The sRouter MUST send the DHCP_RELEASE message [RFC 2131] for the IPv4 public address assigned by DHCPv4 to the sRouter's Operator-Facing Interface whenever the sRouter receives a DHCPv4 server renewal response contains a different IPv4 address. The sRouter MUST NOT wait for a confirmation of the receipt of the release by the DHCPv4 server in order to re-initialize.

7.4 Customer-Facing IPv4 Address Release Behavior

After initiating an administrative device reset in which the public address has been released, the sRouter Customer-Facing interfaces will be limited to inter-LAN forwarding until the device completes any necessary resets and a new address lease is acquired. Prior to the operator-facing interface acquiring an IPv4 address from the operator's DHCPv4 server, local network services and data forwarding of the Customer-Facing LAN interfaces will continue so long as the DHCPv4 server of the sRouter is enabled.

8 OPERATOR-FACING IPV6 PROVISIONING

If the sRouter is operating in either 'IPv6 Protocol Enabled' mode or 'Dual IP Protocol Enabled' mode as defined in Section 6, the sRouter MUST use DHCPv6 [RFC 3315] in order to obtain an IP address for its Operator-Facing IP Interface and any other parameters needed to establish IP connectivity, as illustrated in Figure 5. The sRouter MUST use DHCPv6 prefix delegation [RFC 3633] in order to obtain an IPv6 prefix for the sRouter's Customer-Facing IP Interfaces and any downstream internal routers (IRs), as well as any other parameters needed to establish IPv6 connectivity within the home or office network.

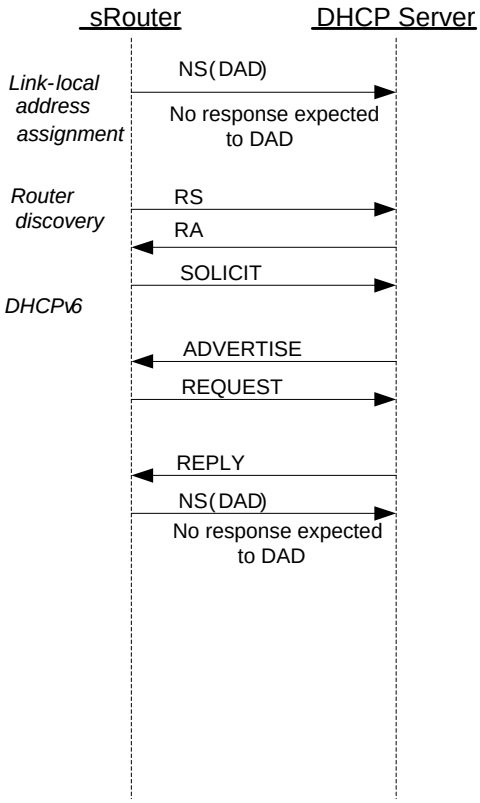

Figure 5 - IPv6 Provisioning Message Flow

The steps in Figure 5 are described in the following subsections.

8.1 Obtain Link-Local Address

The sRouter MUST construct a link-local address for its Operator-Facing Interface and each of its Customer-Facing Interface(s) according to the procedure in section 5.3 of [RFC 4862]. The sRouter MUST use the EUI-64 identifier as a link-local address for each of its interfaces as described in [RFC 4291]. For each of its interfaces, the sRouter MUST join the all-nodes multicast address and the solicited-node multicast address of the corresponding link-local address [RFC 4862], [RFC 2710]. The sRouter MUST use Duplicate Address Detection (DAD), as described in section 5.4 of [RFC 4862], to confirm that the constructed link-local addresses are not already in use prior to sending any sRouter Solicitations on the interface. If the sRouter determines that the constructed link-local address is already in use, the sRouter MUST terminate IPv6 operation on that interface.

8.2 Perform Router Discovery

The sRouter MUST perform router discovery as specified in section 6.3 of [RFC 4861] on its Operator-Facing Interface. The source address used in the sRouter Solicitation MUST be the link-local address on the Operator-Facing Interface. The sRouter identifies neighboring routers and default routers from the received RAs.

8.3 Obtain IPv6 Address and Other Configuration Parameters

The sRouter MUST examine the contents of RAs it receives on the Operator-facing interface and obeys the following rules:

- If the M bit in the RA is set to 1, the sRouter MUST use stateful DHCPv6 to obtain its IA_NA IPv6 address and other configuration information (and ignore the A and O bits).
- If the M bit is set to 1 in the RA, the sRouter MUST use stateful DHCPv6 to obtain its IA_PD.
- If the M bit is set to 0 and the O bit is set to 1, then the sRouter MUST perform stateful DHCPv6 to obtain its IA_NA, IA_PD and other configuration information.
- If both the M bit and the O bit in the RA are set to 0, the sRouter MUST NOT attempt to use DHCPv6 to obtain its IPv6 address and other configuration information.
- The sRouter MUST NOT support SLAAC on its Operator-facing interface.

If the sRouter receives an RA where the M bit is set to zero then the sRouter considers provisioning to have failed.

If an RA contains a prefix advertisement for an IPv6 network prefix on which the sRouter does not have an address and the M bit in the RA is set to 1, the sRouter MUST use DHCPv6 to obtain its IPv6 address for its Operator-Facing Interface and renew any current IA_PD lease(s).

Table 6 depicts sRouter behavior based on the values present in the M and O bits.

Table 6 - sRouter Behavior

M Bit	O Bit	Stateful DHCPv6	Stateless DHCPv6	Prefix Delegation
1	1	Yes	No	Yes
1	0	Yes	No	Yes
0	1	Yes	No	Yes
0	0	No	No	No

The sRouter MUST follow the recommendations in section 4 of [RFC 5942], and in particular the handling of the L flag in the sRouter Advertisement Prefix Information option (PIO).

The sRouter MUST act as a requesting router for the purposes of DHCPv6 prefix delegation ([RFC 3633]). DHCPv6 address assignment (IA_NA) and DHCPv6 prefix delegation (IA_PD) SHOULD be done as a single DHCPv6 session.

The sRouter sends a DHCPv6 Solicit message as described in section 17.1.1 of [RFC 3315]. The Solicit message MUST include:

1. A Client Identifier option containing the DHCP Unique Identifier (DUID) for this sRouter (as specified by [RFC 3315]), the DUID is to be formatted as follows:
   a. The sRouter MUST use a DUID that is one of DUID-LL, DUID-EN or DUID-LLT type and;
   b. The sRouter MUST use a DUID that is persistent across administrative reset or reboot following a loss of power per [RFC 7084] W-6.
2. An IA_NA option to obtain its IPv6 address.
3. An IA_PD option (as specified in [RFC 3633]) to obtain its delegated IPv6 prefix.

4. A Reconfigure Accept option to indicate the sRouter is willing to accept Reconfigure messages.
5. An Options Request option, which MUST include the following options:
    - DNS Recursive Name Server option [RFC 3646].
    - DNS Domain Search List option as per [RFC 3646].
    - OPTION_SOL_MAX_RT (82) as per [RFC 7083].
    - OPTION_NTP_SERVER (56), [RFC 5908].
6. A Vendor-specific option, containing:
    - The 32-bit number 4491 (the Cable Television Laboratories, Inc., enterprise number).
    - A CableLabs Vendor Specific Option Request Option CL_OPTION_ORO as defined in [CANN DHCP-Reg].

The sRouter MUST include the DHCP options for the MAP-E and MAP-T container options in the DHCP Solicit and Request messages per Section 9.7.2.1.

The sRouter MUST use the delegated prefix assigned by the most recent DHCPv6 operation even if the new prefix differs from a prefix that was previously assigned. This new prefix will overwrite any stored prefix information preserved across resets by the sRouter.

If the sRouter does not have a previously assigned delegated prefix, the sRouter MUST indicate a non-zero prefix size as DHCPv6 "hint" information [RFC 3633]. The sRouter MUST ask for a prefix large enough to assign one /64 for each of its Customer-Facing Logical Interfaces rounded up to the nearest nibble. The sRouter MUST be able to accept a delegated prefix length different from what was provided in the hint. If the delegated prefix is too small to address all of its interfaces, the sRouter SHOULD assign a single /64 for all Customer-Facing Logical Interfaces and log an error message.

Any packet received from the Operator-Facing interface by the sRouter with a destination address in the prefix(es) delegated to the sRouter but not in the set of prefix(es) assigned by the sRouter to the Customer-Facing interface MUST be dropped. For example, if the delegated prefix is a /56 but only 12 /64 are in active use, the sRouter should discard all traffic destined to the 242 unused /64 prefixes. This is necessary to prevent forwarding loops and is also helpful in preventing malicious (DoS, network scanning, etc.) traffic from entering the LAN or using sRouter resources.

The sRouter MUST use the following values for retransmission of the Solicit message (see section 14 of [RFC 3315] for details).

- IRT (Initial Retransmission Time) = SOL_TIMEOUT
- MRT (Maximum Retransmission Time) = SOL_MAX_TIMEOUT
- MRC (Maximum Retransmission Count) = 0
- MRD (Maximum Retransmission Duration) = 0

The sRouter MUST use the following value for the Max Solicit timeout value as per [RFC 7083] in preference to any value shown in [RFC 3315]:

- SOL_MAX_RT = 3600 secs

The DHCP server responds to Solicit messages and Request messages with Advertise and Reply messages. The Advertise and Reply messages may include other configuration parameters, as requested by the sRouter, or as configured by the administrator, to be sent to the sRouter. If any of the following options are absent from the Advertise message, and the SOL_MAX_RT option is not present, the sRouter MUST discard the message and wait for another Advertise message. If any of the following options are absent from the Reply message, the sRouter MUST consider IPv6 provisioning to have failed, discard the Reply, and continue transmitting Solicit messages. In addition, the sRouter MAY log an event:

1. The IA_NA option containing the sRouter's IPv6 address;

2. The IA_PD option containing the delegated IPv6 prefix for use by the sRouter;
3. Reconfigure Accept option;
4. The DNS Recursive Name Server Option;
5. DHCP Option 94 (MAP-E) and DHCP Option 95 (MAP-T) are both present.

When the SOL_MAX_RT option is present in an Advertise message, and one or more critical options from the list above are absent, the sRouter MUST acquire the value of SOL_MAX_RT and use such value for future transmissions of Solicit messages. After the sRouter obtains an Advertise containing a new value for SOL_MAX_RT, but lacking critical options from the list above, it MUST use the newly acquired value of SOL_MAX_RT for any subsequent Solicit transmissions. It is not necessary to preserve the value of SOL_MAX_RT across resets.

The sRouter MAY log an event if IPv6 provisioning has failed.

The sRouter interface MUST join the All-Nodes multicast address and the Solicited-Node multicast address of the IPv6 address acquired through DHCPv6. The sRouter MUST perform Duplicate Address Detection (DAD) with the IPv6 address acquired through DHCPv6.

If the sRouter determines through DAD that the IPv6 address assigned through DHCPv6 is already in use by another device, the sRouter MUST:

- Send a DHCP Decline message to the DHCP server, indicating that it has detected that a duplicate IP address exists on the link.
- Discontinue using the duplicate IP address.
- Consider the IPv6 provisioning process to have failed, log the event in the local log, and re-initiate the DHCP process.

The sRouter MUST support the Reconfigure Key Authentication Protocol, as described in section 21.5 of [RFC 3315].

The sRouter MUST NOT forward any IPv6 traffic between its Customer-Facing Interface and its Operator-Facing Interface until it has successfully completed the IPv6 provisioning process. The sRouter MUST NOT forward any IPv6 traffic between its Customer-Facing Interface and its Operator-Facing Interface if, at any time, it does not have a Globally-assigned IPv6 address on its Operator-Facing Interface. The sRouter MUST NOT forward any IPv6 traffic between its Customer-Facing interface and its Operator-Facing interface if it has not completed the delegated prefix acquisition process.

If DHCPv6 provisioning fails on the Operator-Facing interface for any reason, then the sRouter MUST transmit a Router Advertisement on the Customer-Facing interface with the router lifetime equal to zero.

8.4 Use of T1 and T2 Timers

The sRouter MUST initiate the lease renewal process when timer Router-T1 expires. The sRouter MUST initiate the lease rebinding process when timer Router-T2 expires. Timers Router-T1 and Router-T2 are called T1 and T2, respectively, in the DHCP specifications. If the DHCP server sends a value for Router-T1 to the sRouter in a DHCP message option, the sRouter MUST use that value. If the DHCP server does not send a value for Router-T1, the sRouter MUST set T1 to 0.5 times the duration of the lease [RFC 3315]. If the DHCP server sends a value for Router-T2 to the sRouter in DHCP message options, the sRouter MUST use that value. If the DHCP server does not send a value for Router-T2, the Router MUST set Router-T2 to 0.875 times the duration of the lease [RFC 3315].

8.5 Customer-Facing IPv6 Provisioning of CPE Devices

An sRouter that has no default routers on its Operator-Facing Interface MUST NOT send Router Advertisements to its Customer-Facing Interfaces with Router lifetime values other than zero. If an sRouter is serving as an advertising router (acting as the Default Router for the PD) and subsequently detects loss of connectivity on its Operator-Facing

Interface, it MUST deprecate itself as an IPv6 default router on each of its Customer-Facing Interfaces. The sRouter MUST then transmit one or more Router Advertisement messages with the Router Lifetime field set to zero.

Per [RFC 7084], whenever the sRouter detects loss of connectivity to the access network on the Operator-Facing Interface the sRouter MUST:

- Set both the Router Lifetime and the preferred Lifetime to zero (0) in the Router Advertisement (RA) messages for each Customer-Facing Interface that has been allocated a prefix from the delegated prefix that was provisioned on the sRouter Operator-Facing Interface;
- Transmit one (1) or more Router Advertisement (RA) messages on the Customer-Facing Interfaces that have been allocated prefixes from the delegated prefix that was provisioned on the sRouter Operator-Facing Interface.

The sRouter MAY log an event associated with the change in link-state of the Operator Facing Interface.

The sRouter MUST detect disruption of connectivity to the access network when the access network demarcation device (e.g., CM or ONU) loses its connection with the access network. How the sRouter determines the access network demarcation has lost its connection is implementation specific.

Upon detecting access network connectivity has been restored, the sRouter MUST send a DHCPv6 SOLICIT to the DHCP server by resetting the back off timer to its lowest value. Prompt transmission of DHCPv6 SOLICIT messages is essential to re-establishing local IPv6 networking and to allow the injection of the assigned PD into the CMTS's routing table. This insures rapid recovery after planned or unplanned outage events.

The sRouter MUST divide the delegated prefix acquired from the IA_PD option per Section 8.3 during the provisioning process into several sub-prefixes to be used for its Customer-Facing IP Interfaces and any downstream IRs.

By default, the sRouter MUST divide the delegated prefix based on the provisioned prefix size and the configurable Topology mode (Section 8.3) as follows:

- If the provisioned IA_PD is smaller than a /56 (e.g., a /60) and the Topology mode is set to "favor depth", the sRouter MUST divide the delegated prefix on two (2)-bit boundaries into four (4) sub-prefixes by default.
- If the provisioned IA_PD is smaller than a /56 (e.g., a /60) and the Topology mode is set to "favor width", the sRouter MUST divide the delegated prefix on three (3)-bit boundaries into eight (8) sub-prefixes by default.
- If the provisioned IA_PD is a /56 or larger and the Topology mode is set to "favor depth", the sRouter MUST divide the delegated prefix on three (3)-bit boundaries into eight (8) sub-prefixes by default.
- If the provisioned IA_PD is a /56 or larger and the Topology mode is set to "favor width", the sRouter MUST divide the delegated prefix on four (4)-bit boundaries into sixteen (16) sub-prefixes by default.
- If the provisioned IA_PD is too small to divide in the manner described, the sRouter MUST divide the delegated prefix into as many /64 sub-prefixes as possible and log an error message indicating the fault.

For example, if the sRouter is set to "favor width" receives a /56 IA_PD during the provisioning process, the sRouter will split the /56 delegated prefix into sixteen /60 sub-prefixes for use within the home or office. In another scenario where an sRouter set to "favor depth" receives a /62 IA_PD during the provisioning process, it would split that /62 delegated prefix into four /64 prefixes for use within the home or office network.

The sRouter MAY support other methods of dividing the provisioned IA_PD; any such methods would have to be configured by the MSO or its customer.

The sRouter MUST generate and assign a /64 prefix for each Customer-Facing IP Interface before sub-delegating any prefixes to downstream routers within the home.

The sRouter MUST allocate these /64 interface prefixes starting from the numerically lowest sub-prefix generated from the division of the IA_PD (as described above). If the sub-prefix is too small to address all of the Customer-Facing IP Interfaces, the sRouter MUST allocate additional /64 interface prefixes from the next, numerically consecutive sub-prefix.

The sRouter MAY reserve additional /64 interface-prefixes for Customer-Facing Logical Interfaces that could be enabled in the future.

After all of the sRouter's Customer-Facing IP Interfaces have been assigned a /64 prefix, the sRouter MUST delegate sub-prefixes to directly attached downstream routers starting from the numerically highest sub-prefix and working down in reverse numerical order. The prefix assignment in reverse order allows for the flexibility of having a contiguous Customer-Facing IP Interface prefix assignment for interfaces that may be enabled after the initial prefix assignment. This includes the most common use case of additional SSID interfaces that may be administratively disabled at the time the sRouter initializes that are later enabled.

If there are not enough sub-prefixes remaining to delegate to all downstream routers, the sRouter MUST log an error message indicating the fault when the pool of available /64 prefixes has been exhausted.

For example, if there is an sRouter set to "favor depth" configured with two (2) Customer-Facing IP Interfaces that receives a provisioned prefix of 3900:1234:5678:9ab0::/60, the prefix assignment would be as follows:

- Customer-Facing Logical Interface #1 would be assigned with the prefix: 3900:1234:5678:9ab0::/64
- Customer-Facing Logical Interface #2 would be assigned with the prefix: 3900:1234:5678:9ab1::/64

The sRouter would delegate sub-prefixes to the directly attached downstream routers starting first with the 3900:1234:5678:9abc::/62 sub-prefix, and next with 3900:1234:5678:9ab8::/62 sub-prefix, and so on.

If the provisioned prefix is too small to address all of its interfaces, the sRouter MUST collapse the Customer-Facing IP Interfaces into a single Interface and assign a single /64, logging an error message indicating the fault. For example, if sRouter with eight (8) Customer-Facing (physical) Interfaces receives a single /64 prefix during the provisioning process, the sRouter will be forced to bind all eight (8) interfaces into the lowest numbered, or primary LAN, creating a single flat network and a single Customer-Facing IP Interface, regardless of the existing LAN or VLAN configuration(s).

The sRouter MUST assign an IPv6 address from the /64 prefix allocated for each Customer-Facing IP Interface. The sRouter SHOULD generate each Customer-Facing IP Interface Identifier using the Modified EUI-64 process as described per [RFC 4291]. The Modified EUI-64 IPv6 Interface Identifier is created by converting the IEEE 802 MAC address assigned to each Customer-Facing IP Interface to an EUI-64 formatted 64-bit address, and complementing the U/L bit; then, pre-pending 64 bits of the prefix acquired under IA_PD in Section 8.3 to create the 128-bit IPv6 Interface Identifier address.

This entire process can be illustrated in the following way:

1. The aggregate prefix is acquired per Section 8.3.
2. The sRouter then breaks the aggregate prefix into sub-prefixes, based on the Topology Mode in Table 13.
    a. If the aggregate prefix is not large enough, it is broken into as many /64 sub-prefixes as possible and logs an error message.
3. The first of these sub-prefixes is further broken into /64 interface-prefixes for use one on each of the sRouter's Customer-Facing Logical Interfaces.
    a. If the sub-prefix is too small to number all Customer-Facing Logical Interfaces, the sRouter uses additional sub-prefixes as needed (in numerical order).
    b. If the aggregate prefix is too small to number all Customer-Facing Logical Interfaces, the sRouter collapses them into a single interface, assigns a single /64 to that interface, and logs an error message.
4. Each Customer-Facing IP Interface is assigned an IP address from the corresponding interface-prefix.
5. The remaining sub-prefixes are delegated via DHCPv6 to directly downstream routers as needed, in reverse numerical order.

The sRouter MUST support SLAAC [RFC 4862] on all Customer-Facing Interfaces. This requirement satisfies IP address allocation on the Customer-Facing Interfaces for any host that does not implement a full DHCPv6 client.

The sRouter MUST support a DHCPv6 server [RFC 3315] on all Customer-Facing Interfaces. This requirement provides the Customer-Facing Interface with the ability to allocate IP addresses to hosts that implement a DHCPv6 client.

The sRouter MUST support Delegating sRouter behavior for the IA_PD Option [RFC 3633] on all Customer-Facing Interfaces. This requirement provides the means to delegate sub-prefixes to routers within the customer's network from the aggregate, delegated prefix assigned by the operator to the sRouter.

The sRouter MUST support Neighbor Discovery for IPv6 as defined in [RFC 4861].

The sRouter MUST advertise itself as a router for its delegated prefix(es) using the Route Information Option, as specified in section 2.3 of [RFC 4191].

The sRouter's sRouter Advertisement (RA) transmission period MUST be configurable from 3 seconds to 1800 seconds for each Customer-Facing Logical Interface. This configuration flexibility is necessary to adapt to conditions for which the [RFC 4862] defined default of a 120-second interval is inadequate. For example, when prefixes are changed and timely notification of such change is essential to maintaining network continuity.

The sRouter MUST implement a 30-second sRouter Advertisement (RA) transmission interval by default for each of its Customer-Facing Logical Interfaces. If the prefix information contained in an RA changes, the sRouter MUST immediately generate and transmit an updated RA.

8.5.1 Additional Customer-Facing IP Interfaces Enabled After Initial Provisioning

If an sRouter Customer-Facing IP Interface is enabled after initial provisioning and the initial prefix delegation, the sRouter MUST continue prefix assignment for this interface from the next available lowest numbered /64 prefix available. To illustrate using the same example as above, if an additional Customer-Facing IP Interface is enabled after the initial prefix assignment, the sRouter would assign this interface with the prefix of 3900:1234:5678:9ab2::/64.

When the sRouter has used all of its sub-prefixes for any reason, the sRouter MUST NOT enable any new Customer-Facing IP Interfaces. When an attempt to enable a Customer-Facing IP Interface fails because there are no available prefixes, the sRouter MUST log an error message indicating the fault.

8.5.2 SLAAC Requirements for sRouter

SLAAC is required for hosts that do not implement a DHCPv6 client.

The /64 prefix length is required for the dynamic numbering of CPE devices using SLAAC [RFC 4862]. The sRouter MUST generate sRouter Advertisements (RA) on each Customer-Facing Interface as per [RFC 4862].

The sRouter MUST include the following in its RA by default:

- A Prefix Information Option with a prefix derived from the prefix acquired under IA_PD in Section 8.3 and both the ICMPv6 options 'flags' L-Bit (On-link) bit and A-Bit (Autonomous) bit set to 1,
- Preferred lifetimes in the Prefix Information Option set equal to or less than the preferred lifetime communicated in the IA_PD option received on the Operator-Facing Interface. This requirement ensures prefix lifetime synchronization between the sRouter aggregate prefix and the prefix/es assigned to each Customer-Facing Interface.

The above L and A settings in the RA will cause CPE devices to use auto-configuration by default for assigning their global IPv6 address.

Once the sRouter has completed Operator-facing DHCPv6 provisioning, it MUST:

- include DNS configuration option RDNSS in its RA messages as specified in [RFC 6106].
- include DNS configuration option DNSSL in its RA messages as specified in [RFC 6106] if OPTION_DOMAIN_LIST (24) option is acquired via Operator-facing DHCPv6 provisioning.
- include the list of DNS servers specified in the OPTION_DNS_SERVERS (23).

- include the list of domain names specified in the OPTION_DOMAIN_LIST (24) option, if acquired via Operator-facing DHCPv6 provisioning.

8.5.2.1 Local Configuration of SLAAC Options

The sRouter MAY provide a mechanism for local configuration of SLAAC for CPE devices. If local configuration is used, the sRouter MUST override the pass through of options received from the cable operator and provide the locally configured options to CPEs.

8.5.3 DHCPv6 Server Requirements for sRouter

The sRouter MUST provide a DHCPv6 server on Customer-Facing Interfaces as described in:

- [RFC 3315] Dynamic Host Configuration Protocol for IPv6 (DHCPv6).
- [RFC 3736] Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6.

The sRouter DHCPv6 server MUST support providing the following DHCPv6 Options:

- [RFC 3646] DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6).
- [RFC 3633] IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6.
- [RFC 5908] Network Time Protocol (NTP) Configuration Option for DHCPv6.
- [RFC 4242] Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6 (DHCPv6).

The sRouter DHCPv6 server MUST be able to manage at least one IA_NA for each client, and at least one address in each IA_NA.

The sRouter DHCPv6 server MUST be able to manage at least one IA_PD for each client and at least one delegated prefix in each IA_PD. The sub-prefix delegated to the client is derived from the aggregate prefix delegated to the sRouter from the cable operator as described in Section 8.5.

The sRouter DHCPv6 server MUST derive the preferred lifetimes for Customer-Facing IA_NA and IA_PD leases from the preferred lifetime acquired in the IA_PD on the Operator Facing Interface. The sRouter DHCPv6 messages sent on the Customer-Facing interface MUST contain the lifetime value greater than zero and equal to or less than the IA_PD lifetime acquired on the Operator Facing Interface.

IA_NA and IA_PD T1 and T2 values are supplied to the Customer-Facing interface(s) in accordance with [RFC 3315], section 22.4 and [RFC 3633], section 9, respectively.

The sRouter MUST generate sRouter Advertisements (RA) on each Customer-Facing IP Interface as per [RFC 4862]. The sRouter RA MUST include the following by default:

- have the M bit set to 1,
- have the O bit set to 1,
- contain a Prefix Information Option with a prefix derived from the prefix acquired under IA_PD in Section 8.3 and both the ICMPv6 options 'flags' L-Bit (On-link) bit and A-Bit (Autonomous) bit set to 1,
- set the preferred lifetime in the Prefix Information Option equal to the preferred lifetimes communicated in the IA_PD option on the Operator-Facing Interface. This requirement ensures prefix lifetime synchronization between the sRouter aggregate prefix and the prefix(es) assigned to each Customer-Facing Interface.

These settings in the RA will direct CPE devices to use DHCPv6 configuration for assigning their global IPv6 address. In most scenarios, an sRouter would make DHCPv6 services available concurrently with SLAAC in order to supply address and other information to hosts of varying capability. Hosts will be presented with a Router Advertisement that includes the M-bit set to indicate DHCPv6 operation in addition to the A-bit set to indicate SLAAC operation and the O-bit set to support stateless DHCPv6 clients.

NOTE: Recent testing shows operating systems will perform both DHCPv6 and SLAAC for address acquisition when the operating system includes a DHCP client and both methods of address acquisition are made available.

The sRouter MUST be able to pass the following set of options received from the cable operator to the DHCPv6 server for configuration of CPEs.

- DNS Recursive Name Server option as specified in [RFC 3646]
- DNS Domain Search List option as specified in [RFC 3646]

The sRouter MAY relax the requirements on non-volatile storage of assigned addresses and delegated prefixes and MAY glean information about assigned addresses and delegated prefixes from Advertise, Renew, and Rebind messages received from clients.

8.5.3.1 Local Configuration of DHCPv6 Options

The sRouter MAY provide a mechanism for local configuration of DHCPv6 options for CPE devices. If local configuration is used, the sRouter MUST override the pass through of options received from the cable operator and provide the locally configured options to CPEs.

8.5.4 Prefix Changes

An sRouter might receive a replacement prefix from the DHCP server (e.g., during a renewal operation on the Operator-Facing Interface). Due to the global nature of IPv6 addressing of CPEs, the sRouter is to deprecate the previously acquired prefix and allocate addressing from the newly acquired prefix whenever this happens.

The sRouter MUST perform CPE provisioning as per Section 8.5 immediately upon receiving a new prefix.

When an sRouter receives updated information for a currently assigned prefix, the sRouter MUST immediately send Router Advertisements (RAs) with the updated prefix information and IPv6 DHCP RECONFIGURE (type 6, Rebind) on all Customer-Facing Interfaces.

8.6 Operator-Facing IPv6 Address Release Behavior

There are a number of situations in which it is desirable for the sRouter to release its associated IPv6 address leases in order to ensure the integrity of the DHCP database. Examples of such circumstances include situations in which the sRouter needs to be administratively reset (say for configuration change, software update or other reason) or a change to the IPv6 address during DHCPv6 renewal.

The sRouter MUST release its lease information prior to an administratively imposed re-initialization of the access network demarcation device (e.g., CM or ONU) in order to prevent loss of the communications path with the DHCP server. The sRouter MUST NOT wait for confirmation of receipt of the release by the DHCPv6 server in order to re-initialize.

The sRouter MUST send a DHCP_RELEASE message [RFC 3315] for the IPv6 IA_NA and IA_PD assigned by the DHCPv6 server to the sRouter's Operator-Facing Interface for the following events:

- whenever the sRouter is instructed to reset
- whenever the sRouter receives a DHCPv6 Reply message containing a different IPv6 prefix or IPv6 address
- whenever the IA_PD is not renewed for any reason

8.7 Customer-Facing IPv6 Address Release Behavior

After initiating an administrative device reset in which the IA_NA and IA_PD addresses have been released, the sRouter Customer-Facing interfaces will be limited to inter-LAN forwarding until the device completes any necessary resets and new address and prefix leases are acquired.

The sRouter MUST send an ICMPv6 'destination unreachable' message (code 5) for packets forwarded to it that use an address from a prefix that has been deprecated.

After initiating a Reset in which the Operator-Facing Interface's IA_NA and IA_PD addresses have been released, the sRouter MUST declare that it is no longer a Default sRouter by setting the Router Lifetime field to zero in the Router Advertisement.

8.8 CER-ID Requirements

- The sRouter MUST assign the CER-ID value for each of its Customer-Facing IP Interfaces for which an IPv6 prefix has been assigned, by using the corresponding GUA (Global Unique Address) assigned to the Customer-Facing IP Interface.
- The sRouter MUST include the DHCPv6 CL_CER_ID option [CANN DHCP] in Advertise or Reply messages containing an IA_PD.
- If the IPv6 address of the sRouter's Customer-Facing IP Interface that established the CER-ID changes for any reason, the sRouter MUST assign a new value for CER-ID to be included in subsequent DHCPv6 messages.

8.9 Router Interface Addressing Using Link-ID

The sRouter MUST support Link-ID IPv4 address generation as defined below. The sRouter Link-ID feature can be enabled by the operator using TR-069 provisioning methods, see [TR-181]. By default, the Link-ID feature MUST be enabled on the sRouter.

The sRouter MUST NOT persist Link-ID based IPv4 addressing across soft-reset or reboot. If soft-reset or reboot occurs, the sRouter waits for a valid IPv6 PD and provisioning instructions in order to enable Link-ID and generate IPv4 addresses using this algorithm. When sRouter is in dual IP protocol enabled mode and Link-ID is enabled, if there is a temporary loss of connectivity between the Operator-Facing Interface and the CMTS, then the sRouter MUST NOT modify Link-ID based IPv4 addressing. In other words, once Link-ID IPv4 addressing has been generated, the sRouter is expected to maintain this addressing until such time as the IPv6 PD changes or the sRouter is reset.

When operating in 'Dual IP Protocol Enabled' mode and Link-ID is enabled, the sRouter MUST generate a unique /24 prefix for each Customer-Facing IP Interface using the 10.0.0.0/8 aggregate prefix and the Link ID generated from the appropriate IPv6 prefix assigned to the Customer-Facing IP Interface.

A unique IPv4 prefix is created using two steps:

1. Use the decimal value 10 for the first octet,

2. Convert IPv6 Link octets to their decimal equivalents for IPv4 octets 2 and 3.

Step #2 is explained in both the following text and diagram. For example, if an sRouter assigns IPv6 prefix 2001:db8:1234:5601::/64 to a Customer-Facing IP Interface, the Link ID for that interface will be hex 5601(bits 49-64 of the IPv6 prefix). The sRouter will use this Link ID to construct the second and third octets of its /24 IPv4 prefix. The second octet for this example is decimal 86 (equivalent of 0x56, the first octet of the Link ID), and the third octet is decimal 1 (equivalent of 0x01, the second octet of the Link ID). Thus, in this example, the sRouter will assign an IPv4 prefix of 10.86.0.0/24 to the Customer-Facing IP Interface. These requirements enable native routing of IPv4 without the need for a routing protocol or NAPT when the sRouter is operating in Dual IP Protocol Enabled mode. Refer to Figure 6 below and Appendix I for further details.

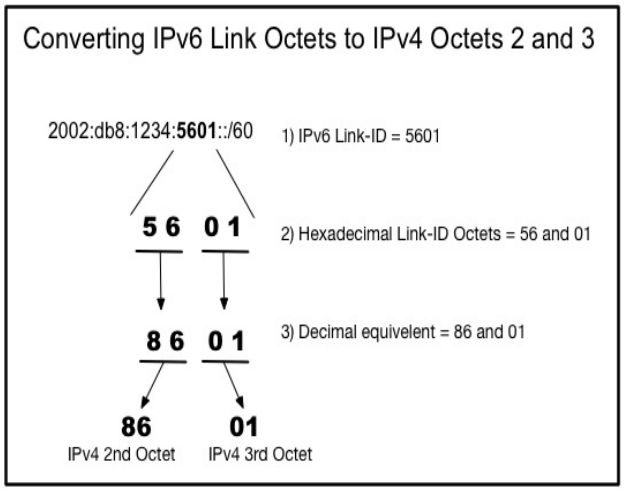

Figure 6 - Example Deriving IPv4 Octets 2 and 3 from an IPv6 Prefix

Using the above Link-ID example, the IPv6 address of 2001:db8:1234:5601::1/64 results in a corresponding IPv4 address of 10.86.1.1/24.

When operating in IPv4 Protocol Enabled mode or when operating in 'Dual IP Protocol Enabled' mode with Link-ID disabled, the sRouter MUST generate each unique /24 IPv4 prefix from one of the three blocks of address space reserved for private internets per [RFC 1918].

9 IPV4 DATA FORWARDING AND NAPT OPERATION

9.1 Introduction

The normative requirements of this section are mandatory for an sRouter that implements the 'IPv4 Protocol Enabled' mode and/or the 'Dual IP Protocol Enabled' mode.

9.1.1 Assumptions

- There is only a single Operator-Facing IP Interface on the sRouter.
- At least one globally-routable IPv4 address is available to the sRouter's Operator-Facing IP Interface.
- The Operator-Facing IP Interface is Ethernet encapsulated.
- The Customer-Facing IP Interface is Ethernet encapsulated.

9.1.2 Overview

IPv4 Forwarding in the sRouter consists of three logical sub-elements:

- IPv4 Router
- NAPT (Network Address Port Translation)
- ARP (Address Resolution Protocol)

The IPv4 Router sub-element is responsible for forwarding packets between the Operator-Facing IP Interface and the Customer-Facing IP Interfaces. This includes looking up the IPv4 Destination address to make a forwarding decision on whether to forward the packet from one of its interfaces to another of its interface or to its internal stack.

Packet handling in the sRouter for NAPT includes:

- Providing a form of IPv4 address translation that allows for multiple IPv4 hosts on the Customer-Facing IP Interfaces while presenting a small number of IPv4 addresses on the Operator-Facing IP Interface.
- Preventing unnecessary traffic on the Customer-Facing IP Interfaces.
- Preventing traffic from one CPE device to another CPE device from traversing to the Operator-Facing Interface.

The ARP protocol on the sRouter provides a mechanism for converting IPv4 network addresses to Ethernet MAC addresses on both Customer-Facing IP Interfaces and the Operator-Facing IP Interface.

9.2 System Description

9.2.1 Overview

Some Routers may have multiple customer ports that are connected to the same logical IP router interface. One scenario would be when the sRouter has an 802.11 wireless port and an 802.3 Ethernet port on the single Customer-Facing logical IP interface. The text in this section uses the term "Customer-Facing IP Interface" to refer to a single Customer-Facing logical IP router interface connected to the sRouter that is not necessarily mapped one-to-one with the number of Customer-Facing ports on the sRouter. This text documents the behavior of a single Customer-Facing IP Interface, though it is possible that an sRouter could have multiple Customer-Facing IP Interfaces.

Packets need to be processed by each of the three sub-elements in a very specific order (see Figure 7). The order is different depending on whether packets are received from a Customer-Facing IP Interface or the Operator-Facing IP Interface.

When receiving packets from the Customer-Facing IP Interface, the sRouter first attempts to route the packet through the router sub-element. If the router sub-element forwards the packet to the Operator-Facing Interface, the packet is passed to the NAPT sub-element to see if the packet requires NAPT translation. Once the NAPT sub-element has completed its work, the packet is sent to the ARP sub-element to resolve the IPv4 network address to

Ethernet MAC. Then the packet is encapsulated in an Ethernet header and sent out the operator interface. If the router sub-element forwards the packet back out the Customer-Facing IP Interface (perhaps because the client is on a different private subnet), the packet is sent to the ARP sub-element to resolve the IPv4 network address to Ethernet MAC. Then the packet is encapsulated in an Ethernet header and sent out the appropriate interface. No NAPT processing is necessary for packets routed back out the Customer-Facing IP Interface.

When packets are received from the Operator-Facing Interface, they are immediately sent to the NAPT sub-element to translate the IPv4 network addresses back to addresses within the domain of the router sub-element. Once the NAPT has been performed on the packet, it is then sent to the router sub-element. If the router sub-element forwards the packet to the Customer-Facing IP Interface, it sends the packet to the ARP sub-element to resolve the IPv4 network address to Ethernet MAC, encapsulates the packet in an Ethernet header, and sends the packet out the appropriate interface. If the router sub-element forwards the packet back to the Operator-Facing IP Interface, it is vendor-specific how to deal with the packet. Some implementations may choose to forward the packet back to the operator network; some may choose to drop the packet. Regardless, traffic should not be sent to a given sRouter from the operator network unless it is destined for a subnet known to the Customer-Facing IP Interface.

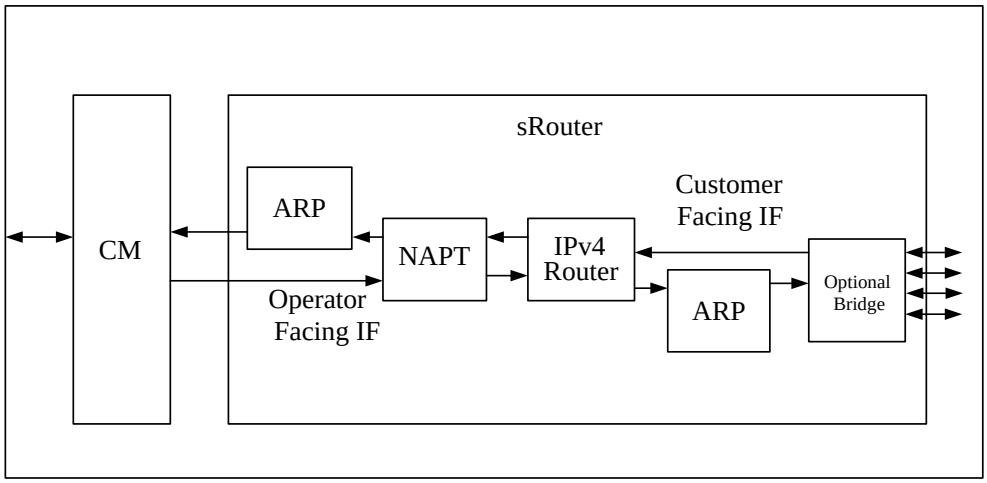

Figure 7 - sRouter IPv4 Forwarding Block Diagram

9.3 IPv4 Router

When the sRouter's IPv4 Router sub-element receives a packet from its NAPT sub-element (received initially by its Operator-Facing IP Interface), it validates the IPv4 header in the packet. The sRouter MAY validate the IPv4 header in accordance with [RFC 1812], section 5.2.2. As defined in [RFC 1812], section 5.3.1, the sRouter MUST decrement the IP TTL field by at least one when forwarding the packet either back to the Customer-Facing IP Interface, or out the Operator-Facing Interface. Packets forwarded to the sRouter's local IP stack for processing MUST NOT decrement the TTL. Once the IPv4 header has been validated, the sRouter processes the destination IPv4 address of the packet. If the destination IPv4 address matches the sRouter's public address assigned to its Operator-Facing IP Interface, the sRouter sends the packet to its local IP stack for processing. If the destination IPv4 address does not match this address, the sRouter determines the next-hop address of the destination in order to forward the packet. The next-hop can be another router or a client directly connected to its Customer-Facing IP Interface. The next-hop is determined by comparing the destination IPv4 address to the subnets assigned to its Customer-Facing IP Interface. If the destination IPv4 address matches any of the prefixes assigned to the Customer-Facing IP Interface, the destination is considered directly connected, or "on-link", and the next-hop to use for ARP purposes is the destination IPv4 address. If it does not match, the destination is considered remote or "not on-link", and the next-hop to use for ARP purposes is the address of the internal router. Discovering other routers on Customer-Facing IP Interfaces, aside from knowledge derived via the use of Link ID when operating in Dual IP Protocol Enabled mode Section 8.9, is vendor-specific. If the sRouter cannot determine the next-hop of the IPv4 destination, then it MUST drop the packet.

When the sRouter's IPv4 Router sub-element receives a packet from its Customer-Facing IP Interface, it validates the IPv4 header in the packet. The sRouter MAY validate the IPv4 header in accordance with [RFC 1812], section 5.2.2. As defined in [RFC 1812], section 5.3.1, the sRouter MUST decrement the IP TTL field by at least one when forwarding the packet, either back to the Customer-Facing IP Interface, or out the Operator-Facing Interface. Packets forwarded to the sRouter's local IP stack for processing MUST NOT decrement the TTL. Once the IPv4 header has been validated, the sRouter processes the destination IPv4 address of the packet. If the destination IPv4 address matches one of the private addresses assigned to the sRouter, it sends the packet to its local IP stack for processing. If the destination IPv4 address does not match one of these addresses, the sRouter determines the next-hop address of the destination in order to forward the packet. The next-hop can be another router or a client directly connected to either its Operator-Facing IP Interface, or back out its Customer-Facing IP Interface. The next-hop is determined by comparing the destination IPv4 address to the subnets assigned to the IP interface on which the sRouter is transmitting. If the destination IPv4 address matches a sub-net prefix, the destination is considered directly connected or "on-link", and the next-hop to use for ARP purposes is the destination IPv4 address. If it does not match, the destination is considered remote or "not on-link", and the next-hop to use for ARP purposes is the address of the intermediate router. The typical scenario for packets routed to the Operator-Facing IP Interface is that the next-hop router will be the sRouter's default, learned via DHCP, Section 7.2, which will be the CMTS. Discovering other routers, aside from the CMTS (or routing delegate chosen by the DHCP server if the CMTS is a bridge) on the Operator-Facing IP Interface, is vendor-specific. Discovery of other directly connected devices on the Operator-Facing IP Interface is also vendor-specific. The typical scenario for packets routed back out the Customer-Facing IP Interface is that the next-hop is a local host on a different subnet than that of the source, but directly connected to the sRouter. If the sRouter cannot determine the next-hop of the IPv4 destination address, it MUST drop the packet.

Regardless of whether the packet was received from the Customer-Facing IP Interface or the Operator IP Interface, the sRouter MUST generate an appropriate ICMP error message as described in [RFC 792] to identify the reason for dropping an IPv4 datagram, except in the follow cases:

- The drop is due to congestion.
- The packet is itself an ICMPv4 error message.
- The packet is destined for an IPv4 broadcast or multicast address.
- The source IPv4 address of the packet is invalid as defined by [RFC 1812], section 5.3.7.
- The packet is a fragment and is not the first fragment (i.e., a packet for which the fragment offset in the IPv4 header is nonzero).

The sRouter's IPv4 router sub-element MUST process and/or generate the following ICMPv4 messages when appropriate:

0	Echo Reply	[RFC 792]
3	Destination Unreachable	[RFC 792]
11	Time Exceeded	[RFC 792]

NOTE: It is considered inappropriate for the sRouter's IPv4 router sub-element to generate ICMPv4 Destination Unreachable messages on the operator-facing interface.

The sRouter MUST have at least one MAC address for its Operator-Facing IP Interface and one MAC address for its Customer-Facing IP Interface. The sRouter MUST share these source MAC addresses for IPv4 and IPv6. The sRouter MUST use the MAC address assigned to its Operator-Facing IP Interface as the source MAC address for all packets that it sends out its Operator-Facing IP Interface. The sRouter MUST use the MAC address assigned to the Customer-Facing IP Interface as the source MAC address for all packets that it sends out its Customer-Facing IP Interfaces.

The sRouter MUST forward broadcast packets received on either interface only to the sRouter's IP stack. The sRouter MUST NOT forward broadcast packets received on either interface to any interface other than the sRouter's IP stack.

9.4 NAPT

The sRouter MUST implement a NAPT function compliant with traditional Network Address Port Translation (NAPT) [RFC 3022], section 2.2. Per [RFC 3022], NAPT "is a method by which many network addresses and their TCP/UDP (Transmission Control Protocol/User Datagram Protocol) ports, are translated into a single network address and its TCP/UDP ports". Also, per [RFC 3022], the purpose of NAPT functionality is to "provide a mechanism to connect a realm with private addresses to an external realm, with globally unique registered addresses". The text in the NAPT sections below uses the term "public address(es)" to refer to the addresses reachable by the sRouter on its Operator-Facing IP Interface, assuming that they are globally-unique registered addresses. Note that an IP address that the sRouter views as globally unique, may be private to the operator's network. However, from the sRouter's perspective, these addresses are unique enough to ensure proper delivery to the next router upstream, and assumed to be globally unique.

Traditional NAPT is the simplest and most straightforward version of NAPT. Other versions that allow for mixtures of public and private network addresses on the Customer-Facing IP Interface, or that allow users from the Operator-Facing IP Interface to establish translations to the Customer-Facing IP Interface, are not required by the sRouter and not discussed in this specification. Traditional NAPT requires that addresses used within the private network on Customer-Facing IP Interfaces cannot overlap with any public addresses reachable by the Operator-Facing IP Interface. Therefore, the sRouter MUST use any of the private IPv4 network addresses described in [RFC 1918] for its Customer-Facing IP Interface.

The sRouter MUST create NAPT translations dynamically based on receiving a packet from a private source on the Customer-Facing IP Interface attempting to access a public address on the Operator-Facing IP Interface, as described in Section 9.4.1.

For packets that traverse the NAPT function, the sRouter MUST always map a combination of private IPv4 address and port number to the same combination of public IPv4 address and port number. That is, the sRouter does not implement a symmetric Network Address Translation (NAT) as defined in [RFC 5389].

The sRouter MUST NOT create NAPT translations when public sources on the Operator-Facing IP Interface attempt to access private destinations on the Customer-Facing IP Interface. Connectivity between two devices that both live on the Customer-Facing IP Interface, but on different subnets, do not require NAPT translations. Therefore, the sRouter MUST NOT create NAPT translations to allow connectivity between CPEs that live on the Customer-Facing IP Interface.

In the following sections, the term Private Network Address Port (PNAP) refers to the network address and TCP/UDP port of a device on Customer-Facing IP Interface that is using a private network address. The term Global Network Address Port (GNAP) refers to the network address and TCP/UDP port of that same device on Operator-Facing IP Interface after it has been translated by NAPT.

9.4.1 Dynamically Triggered NAPT Translations

Dynamically-triggered NAPT is invoked when a device on the Customer-Facing IP Interface with a private network address attempts to initiate one or more sessions to a public destination on the Operator-Facing IP Interface. In this case, the sRouter creates a mapping of source PNAP to GNAP and simultaneously creates a mapping of destination GNAP to PNAP for the return packets. The sRouter then replaces the source PNAP fields of the packet with its corresponding GNAP fields and forwards the packet out the Operator-Facing IP Interface. Once the external destination responds, the sRouter intercepts the reply and changes the previously inserted GNAP fields (now destination) back to the original PNAP values.

The sRouter MUST timeout dynamically-created NAPT translations to ensure that stale entries get removed. The sRouter MUST set the dynamically-created NAPT translations timeout default value to 300 seconds. The sRouter MAY support a configurable time-out value for dynamically-created NAPT translations. Other mechanisms can be used (like analyzing TCP session state) to time out the translations sooner, but the sRouter MUST still time out translations based on the timeout time in case the more advanced mechanism fails (e.g., because packet loss occurred and the sRouter did not see the final packets of a TCP flow).

9.4.2 Application Layer Gateways (ALGs)

Many applications are hampered by NAPT for various reasons. A common problem is the appearance of IPv4 address and/or port information inside the application payload that is too deep into the packet to be manipulated by NAPT, which operates at the network and transport layers. ALGs can be deployed to work around some of the problems encountered, but if the payload of such packets is secured, (by secure transport or application level security) the application cannot work. Another common reason NAPT causes problems is when applications exchange address/port information to establish new connections, creating interdependencies that NAPT cannot know about. The subsections following describe specific ALGs required by the sRouter.

9.4.2.1 ICMP Error Message ALG

ICMP error messages are required for the well-known trace-route network debugging tool to work across the sRouter. This ALG is described in detail in [RFC 3022], section 4.3. The ICMP error message ALG MUST be implemented by the sRouter. Briefly stated, the sRouter MUST translate both the outer and inner IPv4 headers in the ICMP error message in order for the protocol to work correctly, when packets traverse through the NAPT sub-element.

9.4.2.2 FTP ALG

FTP is a fairly widely-used protocol, so the FTP ALG is one of the most important ALGs. The issue with FTP is that it uses the body of the control session packets to signal the data session parameters, including the new TCP ports, to use for the data session. Since NAPT relies heavily on the TCP port field in order to translate between the private and public realm, this ALG is necessary to understand the new ports to be used by the ensuing data session. This ALG is described in detail in [RFC 3022], section 4.4. The FTP ALG MUST be implemented by the sRouter.

9.5 ARP

The ARP function in the sRouter is defined and MUST be compliant with the following RFCs:

- An Ethernet Address Resolution Protocol [RFC 826].
- Requirements for IP Version 4 Routers [RFC 1812], section 3.3.2.
- Requirements for Internet Hosts [RFC 1122], section 2.3.2.

If the corresponding IPv4 network address is found in the table, its corresponding Ethernet address MUST be used as the Ethernet destination address of the packet. If the corresponding IPv4 network address is not found, the sRouter MUST start the ARP protocol in hopes that it will learn the IPv4 network address to Ethernet address association. The sRouter MUST use its own MAC address, as described in Section 9.3, as the source MAC address and source hardware address of all ARP packets.

The sRouter dynamically creates ARP translations based on receiving ARP requests and/or replies for any of its IPv4 network addresses.

ARP entries maintained by the sRouter need careful examination before being aged out of the table. Both voice and video applications present negative subjectively noticeable impacts when ARP entries are removed during a session. [RFC 1122] suggests several different ways to age ARP entries in section 2.3.2.1. The sRouter SHOULD use option 2 – "Unicast polling", which allows for the ARP entry to stay fresh and in the ARP table as long as possible. This option is well-suited for routers that expect to have fairly small ARP tables and want long-term uninterrupted connectivity.

9.6 IPv4 Multicast

The sRouter learns IP multicast group membership information received on the Customer-Facing Interfaces and proxies it on the Operator-Facing Interface towards the next upstream multicast router. The sRouter forwards IPv4 multicast packets downstream based on the information learned at each Customer-Facing Interface.

The sRouter proxies IGMP information upstream actively by implementing mutually-independent IGMPv3 router functionality on Customer-Facing Interfaces, and IGMPv3 group member functionality on the Operator-Facing

Interface. On each IP interface, and independently of other IP interfaces, the sRouter generates, terminates, and processes IGMP messages according to IGMPv3 requirements. For example, the version of IGMP used on the cable network or the local area network will be defined locally at each network. The sRouter may send IGMPv2 reports on the Operator-Facing Interface while generating IGMPv3 queries on Customer-Facing Interfaces.

The following elements define the sRouter IPv4 multicast behavior (also shown in Figure 8):

- An IGMPv3 Group Member that implements the group member part of IGMPv3 [RFC 3376] on the Operator-Facing Interface.
- An IGMPv3 Router that implements the router portion of IGMPv3 [RFC 3376] on each Customer-Facing Interface.
- A subscription database per Customer-Facing Interface with multicast reception state of connected CPEs.
- An IPv4 Group Membership Database that merges subscription information from all the Customer-Facing Interfaces.

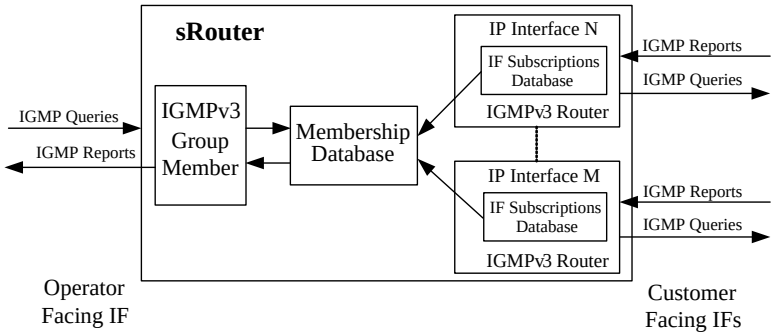

Figure 8 - sRouter IPv4 Multicast Forwarding Block Diagram

Central to the operation of the IGMPv3 Router(s) and IGMPv3 Group Member is the IPv4 Group Membership Database, through which the IGMPv3 Router(s) and IGMPv3 Group Member indirectly relate. This database condenses multicast reception state collected by the IGMPv3 Router(s) from connected CPEs. This information is used by the IGMPv3 Group Member on the Operator-Facing Interface as its own multicast reception interface state.

9.6.1 IGMP Proxying

The sRouter maintains the multicast reception state of CPEs on each Customer-Facing Interface in the interface's multicast subscription database. The sRouter obtains multicast reception state information of CPEs through the implementation of an IGMPv3 Router on each Customer-Facing Interface. Multicast reception state arrives at the sRouter in the form of IGMP Report messages transmitted by CPEs. The sRouter MUST implement the router portion of IGMPv3 [RFC 3376] on each Customer-Facing Interface. The sRouter MUST maintain, for each Customer-Facing Interface, the IPv4 multicast reception state of connected CPEs.

In the event of multiple queriers on one subnet, IGMPv3 elects a single querier-based on the querier IP address. However, the querier election rules defined for IGMPv3 do not apply to the sRouter. The sRouter MUST always act as an IGMP querier on its Customer-Facing Interfaces.

On the Operator-Facing Interface, the sRouter MUST implement the group member portion of IGMPv3 [RFC 3376]. The sRouter MUST merge the multicast reception state of connected CPEs into an IPv4 group membership database as described in Section 9.6.1.1. The sRouter MUST use the IPv4 group membership database as multicast reception interface state per [RFC 3376], section 3.2, on the Operator-Facing Interface. Thus, when the composition of the group membership database changes, the sRouter reports the change with an unsolicited report sent on the Operator-Facing Interface. When queried by an upstream multicast router, the sRouter also responds with information from the group membership database.

The sRouter MUST NOT perform the router portion of IGMPv3 on the Operator-Facing Interface.

9.6.1.1 IPv4 Group Membership Database

The sRouter's Membership Database is formed by merging the multicast reception state records of Customer-Facing Interfaces. In compliance with [RFC 3376], the sRouter keeps per Customer-Facing Interface and per multicast address joined one record of the form:

(multicast address, group timer, filter-mode, (source records))

With source records of the form:

(source address, source timer)

The sRouter keeps an IPv4 Group Membership Database with records of the form:

(multicast-address, filter-mode, source-list)

The sRouter uses the IPv4 Group Membership Database records as the interface state for the IGMPv3 Group Member implementation on the Operator-Facing Interface. Each record of the IPv4 Group Membership Database is the result of merging all subscriptions for that record's multicast-address on Customer-Facing Interfaces. For each IPv4 multicast group joined on any Customer-Facing Interface, the sRouter MUST abide by the following process to merge all customer interface records for the group, into one Group Membership Database record:

- First, the sRouter pre-processes all customer interface group records by:

  Converting IGMPv1 and IGMPv2 records into IGMPv3 records.

  Removing group and source timers from IGMPv3 and converted records.

  Removing every source whose source timer is greater than zero from records with a filter mode value of EXCLUDE.

- Then the sRouter creates an IPv4 Group Membership Database record by merging the pre-processed records, using the merging rules for multiple memberships on a single interface specified in section 3.2 of the IGMPv3 specification [RFC 3376].

9.6.1.2 IPv4 Multicast Forwarding

The forwarding of IPv4 multicast packets received on any interface onto a Customer-Facing Interface is determined by the known multicast reception state of the CPEs connected to the Customer-Facing Interface. The sRouter MUST replicate an IPv4 multicast session on a Customer-Facing Interface, if at least one CPE device connected to the interface has joined the session. The sRouter MUST NOT replicate an IPv4 multicast session on a Customer-Facing Interface if no CPE device connected to the interface has joined the session.

The sRouter MUST NOT forward IPv4 multicast packets received on any interface (i.e., any Customer-Facing or the Operator-Facing Interface) back to the same interface.

The sRouter MUST NOT forward IGMP messages received on any IP interface onto another IP interface.

The sRouter MUST forward IPv4 Local Scope multicast packets (239.255.0.0/16) to all Customer-Facing Interfaces within the same Customer-Facing IP Interface except the Customer-Facing Interface from which they were received.

Except for IGMP packets and IPv4 administratively scoped (239.0.0.0/8) packets, the sRouter MUST forward all IPv4 multicast traffic received on Customer-Facing Interfaces onto the Operator-Facing Interface.

9.6.1.2.1 IPv4 Multicast Forwarding Example

The sRouter in this example has two Customer-Facing Interfaces: CFIA and CFIB, connected to one LAN segment each. On CFIA, there are two CPEs connected: CPE1 and CPE2. CPE1 is IGMPv2 capable and will attempt to join group 224.0.100.1. CPE2 is IGMPv3 capable and will attempt to join group 224.128.100.1 from all sources. On CFIB, there is one CPE connected, CPE3, which is IGMPv3 capable and that will attempt to join group 224.128.100.1, except from source 198.200.200.200.

The router upstream of the sRouter (e.g., the CMTS) supports and is configured to operate in IGMPv3 mode, and thus the sRouter works in IGMPv3 mode on the Operator-Facing Interface.

The setup is shown in Figure 9:

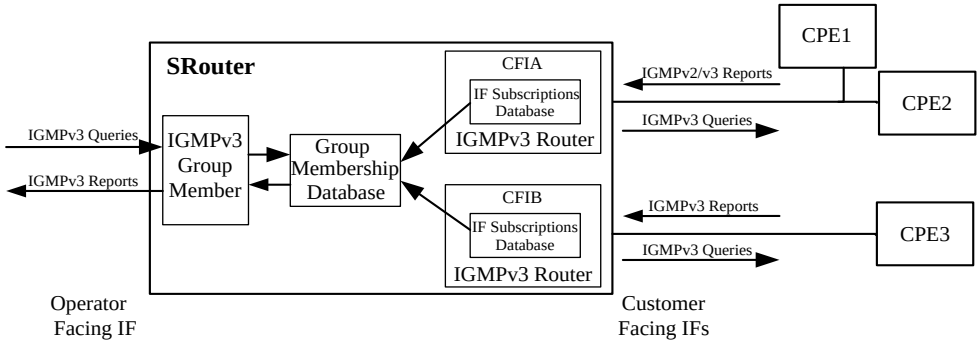

Figure 9 - IPv4 Multicast Forwarding Example

The CPEs send reports as follows:

Report From	Report Version	Multicast Address	Record Type	Source Address
CPE1	IGMPv2	224.0.100.1	N/A	N/A
CPE2	IGMPv3	224.128.100.1	EXCLUDE	Null
CPE3	IGMPv3	224.128.100.1	EXCLUDE	198.200.200.200

Because CPE1 sends an IGMPv2 report for group 224.0.100.1, CFIA operates in IGMPv2 compatibility mode for this group. On the other hand, CFIA and CFIB operate in IGMPv3 mode for group 224.128.100.1, because they receive IGMPv3 reports for this group from CPE2 and CPE3, respectively. The sRouter multicast reception state at each Customer-Facing Interface is the following:

Interface	Multicast Address	Group Timer	Filter-Mode	Source Address	Source Timer
CFIA	224.0.100.1	A	EXCLUDE	Null	0
CFIA	224.128.100.1	B	EXCLUDE	Null	0
CFIB	224.128.100.1	C	EXCLUDE	198.200.200.200	0

The interface state at the sRouter's Operator-Facing Interface, stored in the Group Membership Database, is the following:

Multicast Address	Filter-Mode	Source Address
224.0.100.1	EXCLUDE	Null
224.128.100.1	EXCLUDE	Null

The sRouter uses the information in the table above as multicast reception state at the Operator-Facing Interface. For example, in response to an IGMPv3 general query, the sRouter sends an IGMPv3 report for the two records shown.

Assuming that the CMTS is transmitting downstream four multicast streams, the sRouter forwards them as follows:

Stream #	Multicast Address	Source Address	sRouter forwards on interfaces	
			CFIA	CFIB
1	224.0.200.2	198.100.100.100	NO	NO
2	224.0.100.1	198.100.100.100	YES	NO
3	224.128.100.1	198.100.100.100	YES	YES
4	224.128.100.1	198.200.200.200	YES	NO

9.7 IPv4/IPv6 Co-existence Technologies

Even as operators migrate customers from IPv4 to IPv6 addressing, and deploy IPv6 more widely in their networks, a significant percentage of Internet resources and content will remain accessible only through IPv4. As a consequence of the slow transition to IPv6 on the part of content providers or CE products, operators require mechanisms to allow customers to continue to access content and resources using IPv4 even after the last IPv4 allocations have been fully depleted. This necessitates multiplexing specific groups of subscribers behind a single IPv4 address, or encapsulating or translating IPv4 into IPv6. This section describes several technologies that solve the problem of IPv4/IPv6 co-existence for service providers.

9.7.1 Dual-Stack Lite Operation

Dual-Stack Lite enables an operator to share IPv4 addresses among multiple customers by combining two well-known technologies: IP in IP (IPv4-in-IPv6) tunneling and NAT. More specifically, Dual-Stack Lite encapsulates IPv4 traffic inside an IPv6 tunnel and sends it to an operator NAT device.

When Dual-Stack Lite is enabled, the sRouter acquires an IPv6 address on its Operator-Facing Interface and learns the address of the operator NAT device via DHCPv6. It encapsulates IPv4 traffic inside IPv6 sourced from its Operator-Facing Interface and destined for the operator NAT device.

To facilitate IPv4 extension over an IPv6 network, the sRouter MAY support Dual-Stack Lite.

If the sRouter supports Dual-Stack Lite, it MUST support Dual-Stack Lite B4 functionality as specified in section 5 of [RFC 6333] with the exception of section 5.3.

Requirements in section 5.3 of [RFC 6333] are replaced by the following requirements:

The provisioning of DS Lite MUST be according to [RFC 6334], and request option code 64 (OPTION_AFTR_NAME) for the AFTR tunnel FQDN endpoint name.

9.7.2 Mapping of Address and Port (MAP)

Mapping of Address and Port (MAP) provides a mechanism for IPv4 network domains to communicate with IPv4 network domains over an IPv6-only network. This is particularly useful for operators that have made significant progress in deploying IPv6 in their networks but are challenged in supporting IPv4-only devices within the subscriber's home network.

An operator can use MAP to share IPv4 addresses among multiple customers or operate on a many-to-one or one-to-one basis. MAP Border Relays interpret a defined sequence of bits in the customer's assigned IPv6 prefix, the Embedded Address (EA), to support stateless operation.

MAP defines two types of transport modes: MAP-E and MAP-T. MAP-E uses [RFC2473] encapsulation as a mechanism for converting the IPv4 packet within an IPv6 header. MAP-T uses stateless translation as defined by [RFC6145] to translate the IPv4 header into an IPv6 header.

The sRouter MUST support MAP-E as defined in [MAP-E]. The sRouter MUST support MAP-T as defined in [MAP-T].

The sRouter MUST support configuration of MAP-E or MAP-T functionality via DHCP options as defined in [RFC 7599]. The sRouter is not required to support configuration of both MAP-E and MAP-T simultaneously. In a typical MAP deployment scenario, a MAP CE installs an IPv4 Default Route that directs non-local traffic through an IPv6 encapsulation or translation process so it may be forwarded on to a MAP BR. The resulting forwarding behavior follows a hub and spoke model, where a MAP CE will send all Default Route matching IPv4 destinations through the BR. In this mode of operation, MAP traffic between MAP CEs that belong to the same MAP domain must traverse the BR. In mesh mode, traffic may be forwarded between MAP CEs without an intervening BR.

The sRouter MUST support mesh mode operation between MAP CEs.

- The sRouter MUST support the use of a Basic Mapping Rule (BMR) as a Forwarding Mapping Rule (FMR).
- The sRouter SHOULD support the explicit provisioning of FMR.

If the F-flag in an S46 Rule option is set, the sRouter MUST enable mesh mode for the applicable BMR.

9.7.2.1 MAP-E or MAP-T Configuration Via DHCP

The sRouter that provisions MAP-E or MAP-T through DHCPv6 option encodings MUST issue one (1) Option Request Option (ORO) (option 6) with the appropriate container option as defined in [RFC 7599]:

- Softwire46 MAP-E Container Option (IANA DHCPv6 option 94) in section 5.1
- Softwire46 MAP-T Container Option (IANA DHCPv6 option 95) in section 5.2

Each MAP transport has particular option codes that are embedded in the applicable container option as defined in [RFC 7599]. These option codes MUST NOT be requested in the DHCPv6 ORO option encoding.

For MAP-E configuration, the sRouter MUST accept the following parameters per [RFC 7599] at a minimum to support MAP-E:

- S46 Rule Option (IANA DHCPv6 option 89)
- S46 BR Option (IANA DHCPv6 option 90)
- S46 Port Parameters Option (IANA option 93)

For MAP-T configuration, the sRouter MUST accept the following minimum parameters per [RFC 7599] in order to support MAP-T:

- S46 Rule Option (IANA DHCPv6 option 89)
- S46 DMR Option (IANA DHCPv6 option 91)
- S46 Port Parameters Option (IANA option 93)

9.7.3 Fragmentation

Packet fragmentation is necessary when an IPv4 packet enters the tunnel and the original packet size exceeds the negotiated tunnel MTU. The original IPv4 packet is handled as follows.

1. If in the original IPv4 packet header the DF (Don't Fragment) flag is SET, the sRouter MUST discard the packet, and it MUST return an ICMP message with type = 3 (unreachable), code = 4 (fragmentation needed and Don’t Fragment was set). The next hop MTU field MUST be set to the size of the tunnel MTU by the sRouter.

2. If in the original IPv4 packet header the DF (Don't Fragment) flag is CLEAR, the sRouter MUST perform fragmentation of any IPv4 packet that will exceed the negotiated tunnel MTU. The sRouter MAY fragment in one of two ways:

    a. Via [RFC 6333] section 5.3 where the original IPv4 packet is encapsulated into the IPv6 payload before fragmentation.

    b. Via alternative method where the original IPv4 packet is fragmented first and then each fragment is placed into a separate IPv6 packet.

3. In the sRouter, the method of fragmentation (a or b) MUST be configurable.

The Router MUST support TCP MSS clamping for IPv4 packets and MUST overwrite the TCP MSS with a value supported by the negotiated tunnel MTU.

10 IPV6 DATA FORWARDING

The normative requirements of this section are mandatory for an sRouter that implements the 'IPv6 Protocol Enabled' Mode and/or the 'Dual IP Protocol Enabled' mode, as defined in Section 6.

Assumptions

- There is only a single Operator-Facing IP Interface on the sRouter.
- There is typically a single Customer-Facing IP Interface on the sRouter.
- The Operator-Facing IP Interface is Ethernet encapsulated.
- The Customer-Facing IP Interface is Ethernet encapsulated.
- The sRouter advertises itself as a router (using ND) on all Customer-Facing Interfaces so clients and routers learn about the sRouter. The sRouter does not send Router Advertisements on its Operator-Facing Interface.

10.1 Overview

The IPv6 sRouter is responsible for implementing IPv6 routing. This includes looking up the IPv6 Destination address to decide which of the sRouter interfaces to send the packet.

The Neighbor Discovery (ND) protocol is required on the sRouter. Like ARP in IPv4, it provides a mechanism for converting IPv6 network addresses to Ethernet MAC addresses on both the Customer-Facing IP Interfaces and the Operator-Facing IP Interface. It also provides a mechanism for the sRouter to advertise its presence, host configuration parameters, routes, and on-link preferences.

Figure 10 shows a block diagram of the IPv6 sRouter with an IPv6 sRouter block and an ND block. The IPv6 functionality, however, does not have the clean separation indicated by these blocks. The IPv6 Routing and Neighbor Discovery blocks are closely intertwined and, therefore, are discussed together under the same subsection.

The IPv6 Router uses a local IPv6 routing table to forward packets. The Router creates the IPv6 routing table upon initialization of the IPv6 portion of the Router and adds entries according to the receipt of Router Advertisement messages containing on-link prefixes and routes.

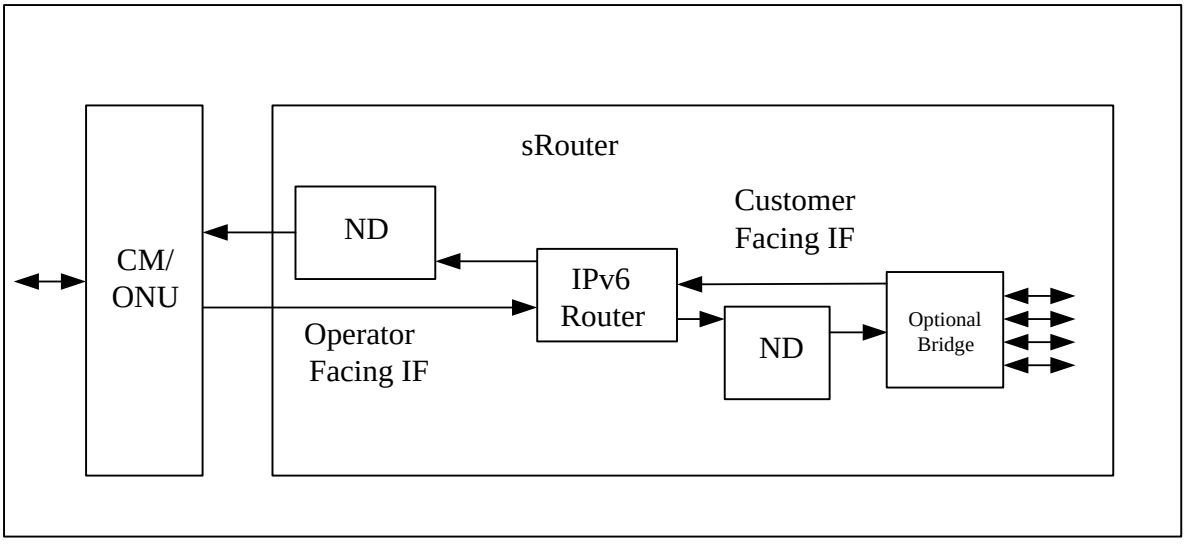

Figure 10 - sRouter IPv6 Forwarding Block Diagram

10.2 System Description

Except when noted, the ND function in the sRouter MUST comply with the Neighbor Discovery for IPv6 [RFC 4861]. Per [RFC 4861], ND is used "to determine the link-layer addresses for neighbors known to reside on attached links and to quickly purge cached values that become invalid."

Several sections of [RFC 4861] do not apply to the SRouter. These sections include:

section 6.2.7 - RA Consistency

section 6.2.8 - Link-local Address Change

section 7.2.8 - Proxy Neighbor Advertisements

section 8 - Redirect Function

section 11 - Security Considerations

section 12 - Renumbering Considerations

The sRouter MUST support the Router Solicitation, Router Advertisement, Neighbor Solicitation, and Neighbor Advertisement messages per [RFC 4861].

The sRouter receives a packet and checks the destination address of the packet. If the destination IPv6 address matches the address assigned to the sRouter's IP interface, the sRouter forwards the packet to its local IP stack for processing. If the destination IPv6 address does not match the sRouter's address, the sRouter determines the next-hop address of the destination in order to forward the packet. The next-hop can be a router, or the destination itself. The next-hop is determined by comparing the destination IPv6 address to the prefixes assigned to the IP interfaces on which the sRouter is communicating. If the destination IPv6 address matches a sub-net prefix, the destination is considered directly connected or "on-link", and the next-hop to use for ND purposes is the destination IPv6 address. If the address of the packet does not match, the destination is considered remote or "not on-link", and the next-hop to use for ND purposes is the address of the intermediate router. If there is no intermediate router, the sRouter MUST immediately drop the packet.

The typical scenario for packets routed to the Operator-Facing IP Interface is that the next-hop router will be the sRouter's default router address, learned via Router Advertisement [RFC 3315], from the service provider’s network. Discovering other routers, aside from the CMTS (or routing delegate if the CMTS/OLT is a bridge), on the Operator-Facing IP Interface is vendor-specific. Discovery of other directly-connected devices on the Operator-Facing IP Interface is also vendor-specific. The typical scenario for packets routed back out the Customer-Facing IP Interface is that the next-hop is a local host on a different subnet than that of the source, but directly connected to the sRouter. If the sRouter cannot determine the next-hop of the IPv6 destination address, then it MUST immediately drop the packet.

Once a next-hop is determined, the sRouter's Neighbor Cache is consulted for the link-layer address of the next-hop address. If necessary, address resolution is performed. Address resolution is accomplished by multicasting a Neighbor Solicitation that prompts the addressed neighbor to return its link-layer address in a Neighbor Advertisement. The neighbor cache entry is then updated with this link-layer address, and the sRouter then forwards the packet to the link-layer address contained in this cache entry. If an error occurs at any point in the process, the sRouter discards the packet. Regardless of whether the packet was received from the Customer-Facing IP Interface or the Operator IP Interface, the sRouter MUST generate an appropriate ICMP error message, as described in [RFC 4884], to identify the reason for dropping an IPv6 datagram, except in the follow cases:

- The drop is due to congestion.
- The packet is itself an ICMPv6 error message.
- The packet is destined for an IPv6 multicast address (except if the packet is the "Packet Too Big Message" or the "Parameter Problem Message", as explained in [RFC 4884], section 2.4, paragraph (e)).
- The packet is destined for a link-layer multicast address.
- The source IPv6 address of the packet does not uniquely identify a single node, as explained in detail in [RFC 4884], section 2.4, paragraph (e).

- The sRouter MUST process and/or generate the following ICMPv6 messages when appropriate:

1	Destination Unreachable	[RFC 4443]
3	Time Exceeded	[RFC 4443]
129	Echo Reply	[RFC 4443]
130	Multicast Listener Query	[RFC 3810]
131	Multicast Listener Report	[RFC 3810]
132	Multicast Listener Done	[RFC 3810]
133	Router Solicitation	[RFC 4861]
134	Router Advertisement	[RFC 4861]
135	Neighbor Solicitation	[RFC 4861]
136	Neighbor Advertisement	[RFC 4861]
143	Version 2 Multicast Listener Report	[RFC 3810]

NOTE: It is considered inappropriate for the Router to generate ICMPv6 Destination Unreachable messages on the operator-facing interface.

The IPv6 CE router MUST implement ICMPv6 according to [RFC 4443].

The sRouter is responsible for decrementing the Hop Limit field in the IPv6 packet that it is going to forward. If the sRouter receives an IPv6 packet with a Hop Limit of zero, or the sRouter decrements an IPv6 packet's Hop Limit to zero, it MUST discard that packet and send an ICMPv6 Time Exceeded message with Code 0 to the source of that IPv6 packet.

The sRouter is also responsible for reinserting the Ethernet header of IPv6 packets. The sRouter has at least one MAC address for its Operator-Facing IP Interface and one MAC address for its Customer-Facing IP Interface that are shared for IPv4 and IPv6 (see Section 8.3). The sRouter MUST use the MAC address assigned to its Operator-Facing IP Interface as the source MAC address for all IPv6 packets that it sends out its Operator-Facing IP Interface. The sRouter MUST use the MAC address assigned to the Customer-Facing IP Interface as the source MAC address for all IPv6 packets that it sends out its Customer-Facing IP Interfaces. Per [RFC 4861], the sRouter uses the MAC address of the next-hop address learned via Neighbor Discovery as the destination MAC address for the IPv6 packet.

The sRouter MUST forward link-local multicast packets received on either interface only to the sRouter's IP stack. The sRouter MUST NOT forward link-local multicast packets received on either interface to any interface other than the sRouter's IP stack.

By default, an sRouter MUST NOT initiate any IPv4 or IPv6 dynamic routing protocols on its Operator-facing interface.

10.3 IPv6 Multicast

The sRouter learns IP multicast group membership information received on the Customer-Facing Interfaces and proxies it on the Operator-Facing Interface towards the next upstream multicast router. The sRouter forwards IPv6 multicast packets downstream based upon the information learned at each Customer-Facing Interface. The sRouter MUST support the forwarding of Any Source Multicast (ASM) and Source-Specific Multicast (SSM) IP multicast streams.

The sRouter proxies MLD information upstream actively by implementing mutually-independent MLDv2 router functionality on Customer-Facing Interfaces and MLDv2 multicast listener functionality on the Operator-Facing Interface. On each IP interface, and independently of other IP interfaces, the sRouter generates, terminates, and processes MLD messages according to MLDv2 requirements. For example, the version of MLD used on the cable network or the local area network will be defined locally at each network. The sRouter may send MLDv1 reports on the Operator-Facing Interface while generating MLDv2 queries on Customer-Facing Interfaces.

The following elements define the sRouter IPv6 multicast behavior (also shown in Figure 11):

- An MLDv2 Multicast Listener that implements the multicast listener part of MLDv2 [RFC 3810] on the Operator-Facing Interface;
- An MLDv2 Router that implements the router part of MLDv2 [RFC 3810] on each Customer-Facing Interface;
- A Subscription Database per Customer-Facing Interface with multicast reception state of connected CPEs;
- An IPv6 Group Membership Database that merges subscription information from all the Customer-Facing Interfaces.

These logical sub-elements are shown in Figure 11.

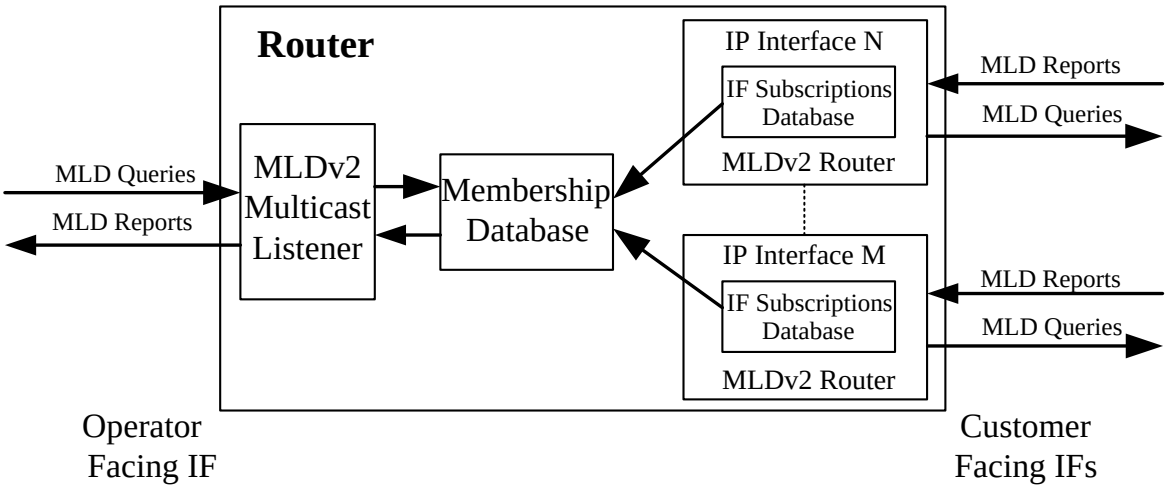

Figure 11 - Router IPv6 Multicast Forwarding Block Diagram

10.3.1 MLD Proxying

The sRouter maintains the multicast reception state of CPEs on each Customer-Facing Interface in the interface's multicast subscription database. The sRouter obtains CPE's multicast reception state information through the implementation of an MLDv2 sRouter on each Customer-Facing interface. Multicast reception state arrives at the sRouter in the form of MLD Report messages transmitted by CPEs. The sRouter MUST implement the router portion of the MLDv2 protocol, [RFC 3810], on each Customer-Facing Interface. The sRouter MUST maintain, for each Customer-Facing Interface, the IPv6 multicast reception state of connected CPEs.

In the event of multiple queriers on one subnet, MLDv2 elects a single querier based on the querier IP address. However, the querier election rules defined for MLDv2 do not apply to the sRouter. The sRouter MUST always act as an MLD querier on its Customer-Facing Interfaces.

On the Operator-Facing Interface, the sRouter MUST implement the multicast listener portion of the MLDv2 protocol, [RFC 3810]. The sRouter MUST merge the multicast reception state of connected CPEs into an IPv6 group membership database, as described in Section 10.3.2, IPv6 Group Membership Database. The sRouter MUST use the membership database as multicast reception interface state per [RFC 3810], section 4.2, for the Operator-Facing Interface. Thus, when the composition of the IPv6 multicast membership database changes, the sRouter reports the change with an unsolicited report sent on the Operator-Facing Interface. When queried by an upstream multicast router, the sRouter also responds with information from the membership database.

The sRouter MUST NOT perform the router portion of MLDv2 on the Operator-Facing Interface.

10.3.2 IPv6 Group Membership Database

The sRouter's Membership Database is formed by merging the multicast reception state records of Customer-Facing Interfaces. In compliance with [RFC 3810], the sRouter keeps per Customer-Facing Interface and per multicast address joined one record of the form:

(multicast address, group timer, filter-mode, (source records))

With source records of the form:

(source address, source timer)

The sRouter keeps an IPv6 Group Membership Database with records of the form:

(multicast-address, filter-mode, source-list)

The sRouter uses the IPv6 Group Membership Database records as interface state for the MLDv2 Multicast Listener implementation on the Operator-Facing Interface. Each record of the IPv6 Group Membership Database is the result of merging all subscriptions for that record's IPv6 multicast-address on Customer-Facing Interfaces. For each IPv6 multicast group joined on any Customer-Facing Interface, the sRouter MUST abide by the following process to merge all customer interface records for the group into one Group Membership Database record:

- First, the sRouter pre-processes all customer interface group records by:

  Converting MLDv1 records into MLDv2 records.

  Removing group and source timers from MLDv2 and converted records.

  Removing every source whose source timer is greater than zero from records with a filter mode value of EXCLUDE.

- Then the sRouter creates an IPv6 Group Membership Database record by merging the pre-processed records, using the merging rules for multiple memberships on a single interface specified in section 4.2 of the MLDv2 specification [RFC 3810].

10.3.3 IPv6 Multicast Forwarding

The forwarding of IPv6 multicast packets received on any interface onto a Customer-Facing Interface is determined by the known multicast reception state of the CPEs connected to the Customer-Facing Interface. The sRouter MUST replicate an IPv6 multicast session on a Customer-Facing Interface if at least one CPE device connected to the interface has joined the session. The sRouter MUST NOT replicate an IPv6 multicast session on a Customer-Facing Interface if no CPE device connected to the interface has joined the session.

In compliance with IPv6 link-scope packet forwarding rules, the sRouter MUST NOT forward MLD messages received on an IP interface onto another IP interface. Also, the sRouter MUST NOT forward link-scoped IPv6 multicast packets received on an IP interface onto another IP interface.

The sRouter MUST forward site-scoped IPv6 multicast packets to all Customer-Facing Interfaces within the same Customer-Facing IP Interface except the Customer-Facing Interface from which they were received.

The sRouter MUST forward all non-link-scoped and non-site-scoped (e.g., not addressed to FF02::/16 or FF05::/16) IPv6 multicast traffic received on Customer-Facing Interfaces onto the Operator-Facing Interface. Operator control of multicast traffic forwarding onto the cable network, if desired, can be done through the implementation of filters at the demarcation device.

10.3.4 IPv6 Multicast Forwarding Example

The sRouter in this example has two Customer-Facing Interfaces: CFIA and CFIB, connected to one LAN segment each. On CFIA, there are two CPEs connected: CPE1 and CPE2. CPE1 is MLDv1-capable and will attempt to join group FF1E::100. CPE2 is MLDv2-capable and will attempt to join group FF1E::128 from all sources. On CFIB, there is one CPE connected, CPE3, which is MLDv2 capable and that will attempt to join group FF1E::128, except from source 3FFE:2900::200.

The router upstream of the sRouter (e.g., the CMTS) supports and is configured to operate in MLDv2 mode, and thus the sRouter works in MLDv2 mode on the Operator-Facing Interface.

The setup is shown in Figure 12:

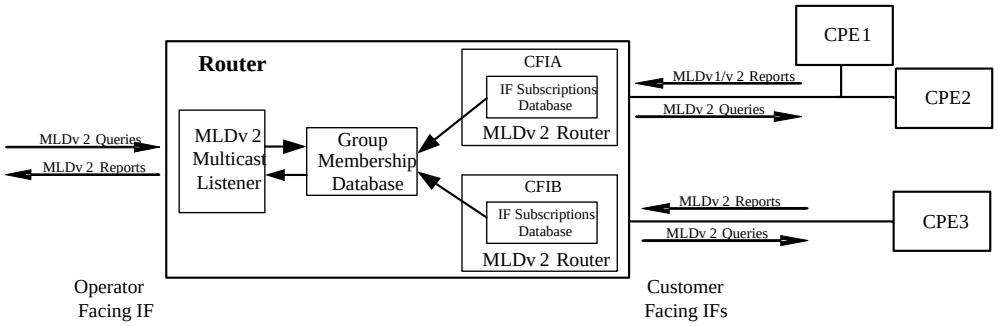

Figure 12 - IPv6 Multicast Forwarding Example

The CPEs send reports as follows:

Report From	Report Version	Multicast Address	Record Type	Source Address
CPE1	MLDv1	FF1E::100	N/A	N/A
CPE2	MLDv2	FF1E::128	EXCLUDE	Null
CPE3	MLDv2	FF1E::128	EXCLUDE	3FFE:2900::200

Because CPE1 sends an MLDv1 report for group FF1E::100, CFIA operates in MLDv1 compatibility mode for this group. On the other hand, CFIA and CFIB operate in MLDv2 mode for group FF1E::128, because they receive MLDv2 reports for this group from CPE2 and CPE3, respectively. The sRouter multicast reception state at each Customer-Facing Interface is the following:

Interface	Multicast Address	Group Timer	Filter-Mode	Source Address	Source Timer
CFIA	FF1E::100	A	EXCLUDE	Null	0
CFIA	FF1E::128	B	EXCLUDE	Null	0
CFIB	FF1E::128	C	EXCLUDE	3FFE:2900::200	0

The sRouter merges the multicast reception state of connected CPEs shown above into the Group Membership Database as follows:

Multicast Address	Filter-Mode	Source Address
FF1E::100	EXCLUDE	Null
FF1E::128	EXCLUDE	Null

The sRouter uses the information in the Group Membership Database as multicast reception state at the Operator-Facing Interface. For example, in response to an MLDv2 general query, the sRouter sends an MLDv2 report for the two records shown.

Assuming that the CMTS is transmitting four multicast streams downstream, the sRouter forwards them as follows:

Stream #	Multicast Address	Source Address	sRouter forwards on interfaces	
			CFIA	CFIB
1	FF1E::200	3FFE:2900::100	NO	NO
2	FF1E::100	3FFE:2900::100	YES	NO
3	FF1E::128	3FFE:2900::100	YES	YES

Stream #	Multicast Address	Source Address	sRouter forwards on interfaces	
			CFIA	CFIB
4	FF1E::128	3FFE:2900::200	YES	NO

11 QUALITY OF SERVICE

QoS on the sRouter is optional. The sRouter SHOULD support Layer 2 and Layer 3 QoS, as defined in this section. The QoS functionality described herein allows the operator to selectively provide a level of differentiation among the various data streams destined for CPE behind the sRouter. Typical applications could include Internet Protocol Television (IPTV) services and other enhanced data services, though it is anticipated that overall packet counts will still be dominated by largely undifferentiated best-effort data traffic.

If the sRouter supports QoS, the sRouter MUST prioritize the forwarding of IP packets based on the values marked in the IPv4 ToS byte or IPv6 Traffic Class field. This is because Layer 2 (e.g., 802.1 p/Q Ethernet) headers will be removed as the packets traverse the sRouter.

11.1 Downstream Quality of Service Operation

This section deals with the requirements regarding traffic going to CPEs, through the sRouter, from the access network.

If the sRouter supports QoS, the sRouter MUST provide two or more priority queues on each Customer-Facing Interface for traffic going to CPEs. The sRouter MAY provide a configuration mechanism to map ToS/Traffic Class field priority values to the high and low priority queues. As a default setting, the sRouter might use the most significant bit of the ToS/Traffic Class field to determine priority to queue mappings.

11.2 Upstream Quality of Service Operation

This section deals with traffic coming from the CPEs attached to the sRouter to the access network.

For the purposes of applying QoS to upstream traffic sourced from CPE devices, the interface between the sRouter and the access network demarcation device is considered to be of infinite bandwidth and thus no congestion, control, priority, nor reservation of bandwidth resources should be expected to occur on this interface. Thus, the sRouter does not need to provide any queues in the upstream direction. The sRouter MAY provide a configuration mechanism to determine whether the sRouter allows CPE devices to pass QoS-tagged packets with the IP ToS/Traffic Class field intact, or whether the sRouter resets the IP ToS/Traffic Class field to 0. The sRouter MAY use the IP ToS/Traffic Class field to populate Layer 2 QoS headers to ensure upstream QoS treatment. Although other implementations are possible, one such implementation is to directly map the three most significant bits of the IP ToS/Traffic Class field into the 802.1Q priority field.

In the case where multiple Customer-Facing Interfaces are implemented, the sRouter may support additional QoS mechanisms to prioritize upstream traffic based on ingress interface.

12 sRouter Management

The sRouter allows the implementation of different management interfaces as described in this section. Management interfaces in this specification refer to the protocols, data models, and semantic representation of the data exchange to perform the conventional management functions in the device.

The sRouter MUST support TR-069 [TR-069] from the Operator-Facing management interface.

The router MAY support SNMP for remote management capability. If the router supports SNMP, then it MUST support SNMPv2 [RFC1901] and [RFC1907].

User management from the Customer-Facing interface is vendor specific. Remote management of the sRouter (from the Operator-Facing interface) by the customer is outside of the scope of this specification.

12.1 sRouter TR-069 Management Interface Requirements

The sRouter TR-069 Management Interface requirements are listed in Annex C.

12.2 sRouter SNMP Management Interface Requirements

The sRouter SNMP management interface is configured using the TR-069 Management Interface. [TR-181] is extended to provision basic SNMP agent functions as listed in [WI-FI MGMT].

12.2.1 ACS Discovery

The eRouter performs initial ACS discovery via the mechanisms in the following sections.

12.2.1.1 TR-069 Management Server DHCP Requirements

The sRouter MUST follow the DHCP requirements in [TR-069] for the initial ACS discovery with the possible exception of using any CableLabs-defined DHCP options mentioned here or elsewhere in this specification.

12.2.2 ACS Selection

The sRouter MUST use TR-069 Management Server URL DHCP Option as the initial ACS URL. If the TR-069 Management Server URL is not present in the DHCP Offer/Response, the sRouter MUST NOT communicate with any ACS.

12.2.3 Dynamic ACS Updates

After the initial discovery, the ACS URL can be changed by updating the Device.ManagementServer URL attribute value. The sRouter MUST ignore the ACS URL if it is present in DHCP renew/rebind messages.

12.2.4 TR-069 CWMP Control and Credentials

The TR-069 Device.ManagementServer object defines controls for CWMP operations and credentials for authentication of connection requests between the CPE and ACS. All TR-069 Device.ManagementServer objects can be configured by the ACS via [TR-069] procedures.

The parameter Device.ManagementServer.URL is delivered via DHCP.

13 SECURITY[1]

sRouter security is important in order to provide a reliable service to subscribers and to protect their traffic as well as the operator's network from attacks. The security requirements in this section cover traffic filtering, the management interface, software update, secure boot, the data plane and community Wi-Fi.

13.1 Traffic Filtering

It is considered a best practice to filter obviously malicious traffic (e.g., spoofed packets, "Martian" addresses, etc.). Thus, the sRouter ought to support basic stateless egress and ingress filters. The sRouter is also expected to offer mechanisms to filter traffic entering the customer network; however, the method by which vendors implement configurable packet filtering is beyond the scope of this document. Filters SHOULD be remotely configurable by the operator. All provided filters MUST be able to be remotely enabled or disabled by the operator. While implementation is left to vendors, filters SHOULD allow protection of both the sRouter and traffic passing through the sRouter.

The sRouter MUST enable a stateful firewall by default. In particular, the sRouter SHOULD support functionality sufficient for implementing the set of recommendations in [RFC 6092], section 4. The sRouter MUST support ingress filtering in accordance with BCP 38 [RFC 2827]. Management and configuration of the stateful firewall may be by either or both the operator or the end user; however, the firewall SHOULD be remotely configurable by the operator. All features of the firewall MUST be able to be remotely enabled or disabled by the operator. Implementation method of the firewall by vendors is beyond the scope of the document.

[RFC 6092] contains 50 "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service." Not all of these recommendations are applicable to MSO networks. Of the applicable recommendations, not all are needed immediately. In order to ensure that vendors are able to implement "simple security" support in sRouter devices, Annex D categorizes the recommendations into five requirement categories:

- Critical - Critical to network connectivity. Include in initial release.
- Important - Failure to implement could open subscribers to infosec attack.
- BCP - Security best practice/nice to have but not critical.
- Other - MSOs have indicated ambivalence to this category of recommendations.
- Conflict - Recommendation conflicts with MSO needs and requires modification or should not be implemented.

The local user SHOULD NOT be able to filter or block traffic necessary for the network operator to manage the sRouter.

13.2 Management Security on the Operator-Facing Interface

The management connection on the operator-facing interface uses TR-069 and may use SNMP. The operator can secure the TR-069 interface by using an HTTPS URL for the Auto-Configuration Server (ACS). The sRouter MUST establish a TLS connection with the ACS when an HTTPS URL is provided in the TR-069 Management Server URL DHCP Option. When TLS is being used, the sRouter MUST NOT accept or respond to TR-069 messages outside of the TLS connection. The sRouter MUST support TLS 1.1 [RFC 4346] and TLS 1.2 [RFC 5246]. The sRouter MAY support other TLS versions. When SNMP is used, the sRouter MUST establish a DTLS connection with the SNMP manager if a secure URI is provided during SNMP management interface configuration. When DTLS is being used, the sRouter MUST NOT accept or respond to SNMP messages outside of the DTLS connection. The sRouter MUST support DTLS 1.0 [RFC 4347] and DTLS 1.2 [RFC 6347]. It may also support other DTLS versions.

When using TLS or DTLS, the sRouter MUST communicate its capabilities to the ACS or SNMP manager as specified in appendix E of TLS 1.2 [RFC 5246], allowing the servers to choose the version of protocol to use.

[1] Modified per sRouter-N-16.0155-2 on 12/7/16 by KB.

During the initial TLS message exchange, the server and client negotiate the security algorithms for key exchange, traffic encryption and traffic integrity using cipher suites. The client indicates what cipher suites it supports and the server selects which cipher suite will be used to secure the TLS connection. The sRouter MUST support the following cipher suites defined in TLS 1.2 [RFC 5246] when using TLS or DTLS:

-TLS_RSA_WITH_AES_128_CBC_SHA256
-TLS_RSA_WITH_AES_256_CBC_SHA256
-TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
-TLS_DHE_RSA_WITH_AES_256_CBC_SHA256

The sRouter may support additional cipher suites.

Digital certificate PKI credentials are used to authenticate both ends of the TLS or DTLS connection. The sRouter MUST use the digital certificate PKI defined in Annex F when establishing a secure TLS or DTLS management connection with the headend.

The root CA certificate of the PKI is used as a trust anchor in both the headend management server and the sRouter for validating received certificates. The root CA certificate can be acquired by requesting it from the CA managing the PKI. The operator installs the root CA certificate as a trust anchor in its management server(s). The sRouter MUST have the root CA certificate installed as a trust anchor for validating certificates received from headend management servers.

A certificate with a unique identifier is installed in the management server(s) and in each sRouter device along with its corresponding private key and issuing intermediate CA certificate. The operator is responsible for acquiring certificates for its management servers and the manufacturer is responsible for acquiring sRouter device certificates to be installed during the manufacturing process. The sRouter MUST have a device certificate, its corresponding private key, and issuing CA certificate installed. When mutual authentication is performed, the server and sRouter exchange certificates (server/device certificate and issuing intermediate CA certificate).

13.2.1 Management Server Authentication

When an sRouter receives certificates from the management server during TLS or DTLS authentication it MUST perform the following validation checks:

1. Verify that the server certificate chains to the root CA certificate trust anchor using Basic Path Validation procedures defined in the X.509 PKI certificate standard [RFC 5280].

2. Verify that the host portion of the ACS URL or SNMP manager URL matches the Common Name attribute of the Subject field or any domain names in the Subject Alternative Name extension in the server certificate.

3. Verify the server certificate and issuing intermediate CA certificate are not revoked by checking the corresponding CRL of the issuing CAs.

To perform certificate revocation the sRouter MUST support CRLs as defined in the X.509 PKI certificate standard [RFC 5280]. During manufacture the sRouter MUST have the latest CRLs from the issuing CA of the server certificate and the intermediate CA certificate installed. When the current date/time is greater than the CRL's nextUpdate time, the sRouter MUST attempt to download a new CRL using the URL provided in the Issuing Distribution Point extension of the CRL. If the attempt to download a new CRL is unsuccessful, the sRouter MUST continue to use the current CRL.

Before accepting a new CRL, the sRouter MUST verify its signature using the CA certificate identified in the current CRL's AIA extension. If the sRouter does not have the CA certificate, it MUST download it using the URL provided in the current CRL's AIA extension. If the sRouter is unable to verify the new CRL's signature, it must reject the new CRL and continue using the current CRL.

If the attempt to download a new CRL or verify a new CRL's signature is unsuccessful, the sRouter MUST log an error, and retry the CRL download using the Session Retry Policy defined in [TR-069]. If the sRouter is unable to establish a successful TLS or DTLS connection with the management server, it MUST retry the connection using the Session Retry Policy defined in [TR-069] and reset itself after the 5th retry count.

13.2.2 sRouter Authentication

To support certificate-based mutual authentication, operators need to configure their management servers to request client certificates during TLS or DTLS certificate exchange. Management servers should validate sRouter device/client certificates using a similar method as defined above. The sRouter MUST send its device certificate and issuing intermediate CA certificate to the management server during the TLS and DTLS authentication process when requested.

13.3 Secure Software Update

The main steps for sRouter secure software update involve downloading software from the file server, installation, and software verification each time it is loaded into memory for execution during boot up. The following sections cover: software download and secure boot.

13.3.1 Secure Software Download

There are two methods used for secure software download: TR-069 and DOCSIS SNMP.

13.3.1.1 TR-069 Secure Download

TR-069 uses HTTP to download a software image from a file server. The download is triggered by the operator sending a Download RPC message to the sRouter. The operator can secure the download connection by using an HTTPS URL in the Download RPC message. If an HTTPS URL is provided in the TR-069 Download RPC message, the sRouter MUST establish a TLS connection with the file server before downloading the software image.

Mutual authentication of the HTTPS download connection uses digital certificate PKI credentials and could use a username/password credential. The certificate PKI is defined in Annex F. The sRouter MUST use the digital certificate PKI defined in Annex F when establishing an HTTPS connection with the download file server and performing server or client certificate authentication.

The root CA certificate of the PKI is used as a trust anchor in both the headend file server and the sRouter for validating received certificates. This is the same root CA certificate as used with securing the management connection in Section 13.2. The root CA certificate can be acquired by requesting it from the CA managing the PKI. The sRouter MUST have the root CA certificate installed as a trust anchor for validating certificates received from the headend file server.

For sRouter authentication a device certificate or username/password credential can be used. If a device certificate is being used, the root CA certificate should be installed on the file server as a trust anchor for validating received sRouter device certificates. If a username/password credential is used, the file server should have access to the authorized username/password so it can be used to verify the username/password received from the sRouter.

A certificate with a unique identifier is installed in the file server along with its corresponding private key and issuing intermediate CA certificate. The operator is responsible for acquiring the certificates for the file server. When a client certificate is used for sRouter authentication, the sRouter device certificate used in securing the management connection can also be used for securing the software download connection (see Section 13.3). To enable sRouter authentication using a device certificate, the operator should configure their file server to request a client certificate when establishing an HTTPS connection. When the sRouter receives a client certificate request during HTTPS authentication, it MUST send its device certificate and issuing intermediate CA certificate to the file download server.

When an sRouter receives certificates from the file server during HTTPS authentication, it MUST perform the following validation checks:

1. Verify that the server certificate chains to the root CA certificate trust anchor using Basic Path Validation procedures defined in the X.509 PKI certificate standard [RFC 5280].

2. Verify that the host portion of the Download RPC URL matches the Common Name attribute of the Subject field or any domain names in the Subject Alternative Name extension in the server certificate.

3. Verify the server certificate and issuing intermediate CA certificate are not revoked. First check using an OCSP request to the OCSP URL in the certificate's AIA extension. If there is no response from the OCSP responder, then the sRouter checks the latest version of the CRL list it has stored. Additional sRouter certificate revocation support requirements are defined in section 13.3.

When device certificates are used for sRouter authentication, the file server should validate them using a similar method as defined above.

After the sRouter has downloaded the new software image, it MUST verify that it is compatible with the sRouter using a vendor-specific method before replacing the current image with the new image.

The image MAY be signed using the method defined in the DOCSIS 3.1 Security specification using code verification certificates (CVCs) issued from the PKI defined in Annex F (same as DOCSIS 3.1 new PKI). If TR-069 image signature verification is enabled via the management interface, the sRouter MUST validate the image signature as defined in the Code Verification Requirements section of the DOCSIS 3.1 Security specification with the PKI defined in Annex F when using the TR-069 download method. If TR-069 image signature verification is disabled and the image has been signed, the sRouter MAY ignore the image signature.

13.3.1.2 SNMP Initiated Secure Software Download

The SNMP initiated secure software download mechanism for the sRouter is based on the DOCSIS secure software download method defined in DOCSIS 3.1 specifications. The software image is signed using a code verification certificate (CVC) and downloaded using TFTP from a file server. The sRouter then validates the signature on the image before replacing the current image. If SNMP management is supported by the sRouter, the DOCSIS SNMP secure software download mechanism defined in this section MUST also be supported unless specified otherwise using the certificate PKI defined in Annex F (the Annex F PKI is the same as the DOCSIS 3.1 new PKI). From a DOCSIS specification perspective, secure software download is considered already enabled and will trigger a download when the appropriate MIBs are set. The configuration file CVC to enable secure software download does not apply.

13.3.2 Secure Boot

To help prevent attackers from hacking the sRouter device and replacing the software image with an unauthorized version, a method to secure the boot image is needed to verify that the image loaded into memory for execution during the boot process is authorized and valid. During the boot process, the sRouter SHOULD verify the software image is authorized before loading it into memory for execution using a hardware protected mechanism; the attacker cannot simply reflash the non-volatile memory in order to bypass the secure boot process. The sRouter SHOULD hardware-protect any keys used to verify the image. Details of the secure boot mechanism are left to vendor implementation detail.

13.4 Physical Protection of Keys in the sRouter

The sRouter MUST store the Device Certificate private key in a manner that deters unauthorized disclosure and modification. Also, the sRouter SHOULD prevent debugger tools from reading the Device Certificate private key in production devices by restricting or blocking physical access to memory containing this key. The sRouter MUST protect trust anchor public keys used for validating certificate chains and image signatures against unauthorized changes.

The sRouter MUST meet [FIPS 140-2] security requirements for all instances of private and public permanent key storage.

The sRouter MUST meet [FIPS 140-2] Security Level 1. [FIPS 140-2] Security Level 1 requires minimal physical protection through the use of production-grade enclosures. The reader should refer to the cited document for the formal requirements; however, below is a summary of those requirements.

Under the [FIPS 140-2] classification of "physical embodiments" of cryptographic modules, sRouter devices are "multiple-chip stand-alone cryptographic modules." [FIPS 140-2] specifies the following Security level 1 requirements for multiple-chip stand-alone modules:

- The chips are to be of production-grade quality, which shall include standard passivation techniques (i.e., a sealing coat over the chip circuitry to protect it against environmental or other physical damage);
- The circuitry within the module is to be implemented as a production grade multiple-chip embodiment (i.e., an IC printed circuit board, a ceramic substrate, etc.);
- The module is to be entirely contained within a metal or hard plastic production-grade enclosure, which may include doors or removable covers.

13.5 Private Data Plane Security on Operator-Facing Interface

The operator-facing data plane carries in-home subscriber Internet traffic between the sRouter and the headend. Traffic may be unsecured (e.g., HTTP) or secured (e.g., HTTPS) by applications embedded in the sRouter or the subscriber's end devices. In cases where there is a reasonable threat of man-in-the-middle attacks, this traffic, or at least the unsecured traffic, should be secured using a protocol such as IPSec since it is applied lower in the protocol stack, including but not limited to transport protocols such as UDP and TCP IP traffic.

13.6 Community Wi-Fi Security

When an sRouter supports community Wi-Fi traffic from roaming subscribers, it may be vulnerable to man-in-the-middle (MitM) attacks. A community Wi-Fi open SSID has no traffic protection over the wireless link, so subscribers should use application-specific security to protect sensitive or confidential information between the mobile device and the application server. This would include securing the connection to the authentication portal using HTTPS. Some MSOs may provide VPN applications to protect traffic.

For community Wi-Fi secure SSIDs, traffic is protected over the wireless link, but a MitM vulnerability can still exist at the modem where bridging of access network security and the Ethernet link occurs. It is important to properly secure community Wi-Fi secure SSID traffic, since roaming subscribers will assume their traffic is protected over the operator's network. This includes RADIUS traffic between the AP and AAA server. After successful authentication, the AAA server sends the pairwise master key (PMK) to the AP for establishing a secure WPA2 connection with the mobile device. This PMK needs to be protected.

If GRE security is enabled via the management interface, the sRouter MUST establish a secure DTLS 1.2 connection with the headend GRE concentrator to secure community Wi-Fi tunneled traffic. Digital certificate PKI credentials are used to authenticate both ends of the DTLS connection. The certificate PKI requirements defined for securing the management interface also apply when securing the GRE concentrator connection. The sRouter MUST use the certificate PKI requirements in Section 13.2 for securing the GRE tunnel, except that it verifies that the host portion of the GRE concentrator URL matches the Common Name attribute of the Subject field or any domain names in the Subject Alternative Name extension in the server certificate. The sRouter device certificate is used for both management and GRE interface client authentication.

If AAA server security is enabled and connection to a AAA server has been configured via the management interface, the sRouter MUST establish a secure DTLS 1.2 connection with the AAA server to protect transmission of the PMK. The certificate PKI requirements defined for securing the management interface also apply when securing the AAA server connection. The sRouter MUST use the certificate PKI requirements in Section 13.2 for securing the AAA server connection, except that it verifies that the host portion of the AAA server URL matches the Common Name attribute of the Subject field or any domain names in the Subject Alternative Name extension in the server certificate. The sRouter device certificate is used for both management and AAA server interface client authentication.

13.7 IPSec

The sRouter MAY utilize an IPSec tunnel or an IPSEC encrypted GRE tunnel to protect traffic on the Operator-Facing Interface. This traffic includes SNMP management traffic as well as AAA management traffic (RADIUS,

TACACS) and GRE traffic for Community Wi-Fi applications. TR-069 management traffic is not included because it is protected by TLS as defined in Section 13.2. The sRouter IPSec architecture consists of an IPSec Gateway located in the SP's headend that terminates IPSec tunnels from sRouter devices. An example architecture is provided in Figure 13.

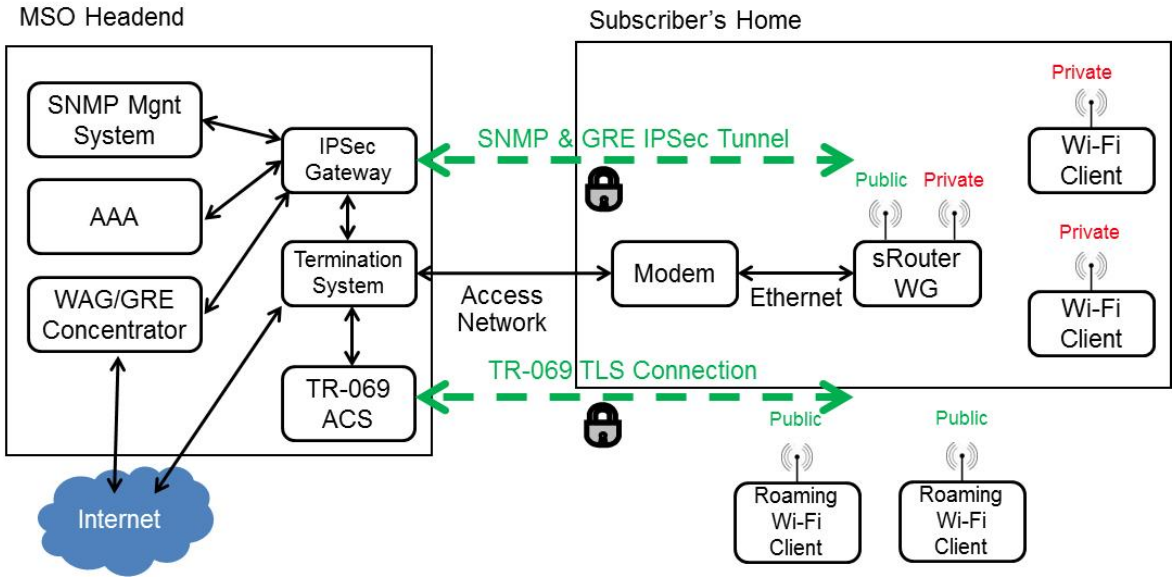

Figure 13 - Example IPSec Architecture

If the sRouter supports IPSec for securing traffic on the operator-facing interface, it MUST implement the following requirements:

- If the sRouter is configured to enable the IPSec tunnel via the TR-069 management interface, it will establish an IPSec tunnel with the IPSec Gateway having the indicated domain name. The sRouter is required to support IPSec as defined in [RFC 4301].
- IKEv2 is used to provide mutual authentication and key exchange between the IPSec Gateway and the sRouter using the certificate PKI defined in Annex F. The sRouter is required to support IKEv2 defined in [RFC 7296]. The sRouter will initiate an IKEv2 message exchange with the IPSec Gateway and use the certificate PKI defined in Annex F for endpoint authentication purposes.
- The sRouter is required to validate certificates received from the IPSec Gateway using the Management Server Authentication steps defined in 13.2.1, except that it verifies the IPSec Gateway domain name value of the TR-069 management configuration matches the Common Name attribute of the Subject field or any domain names in the Subject Alternative Name extension in the gateway certificate.
- During IKEv2 certificate exchange, the sRouter will send its device certificate and issuing intermediate CA certificate to the IPSec Gateway.
- The sRouter will use IPSec tunnel mode with an ESP header to secure packets.
- The sRouter is required to support the following cipher suites:
  - Data Encryption - AES 128 CVC, AES 256 CBC
  - Message Integrity - HMAC-SHA1-96, HMAC-SHA2-256-128
  - Pseudo-random function for key generation - HMAC-SHA1, HMAC-SHA2-256
- The sRouter is required to use the IPSec tunnel for sending & receiving SNMP, AAA, and GRE traffic. The IPSec tunnel can also be used for other traffic, e.g., the private data plane. The sRouter is required to ignore any SNMP, AAA, or GRE traffic received outside of the IPSec tunnel.

- If the sRouter is unable to establish a successful IPSec connection with the IPSec Gateway, it is required to retry the connection using the Session Retry Policy defined in [TR-069] and reset itself after the 5th retry count.

14 Router Tunnel Management and Configuration

14.1 GRE Requirements

Some of the applications envisioned for the sRouter rely upon the tunneling of traffic between the customer's service location and the Operator's network core. For example, a community Wi-Fi application that utilizes one or more SSIDs to provide public Wi-Fi could be configured to tunnel its traffic to a central concentrator within the Operator's network core. While multiple tunneling protocols and techniques exist, Generic Route Encapsulation (GRE) tunneling, see [RFC 2784] and [RFC 2890], has become the prevalent method for conveying traffic to the Operator's core. In order to support the management and configuration of these GRE tunnels, this specification defines both SNMP-based MIBs and TR-181 data model elements. The SNMP MIB defined referenced in this specification originated from the [TR-181] data model profiles for GRE tunnels.

An sRouter that implements GRE over IPv4 MUST support [RFC 2784].

An sRouter that implements GRE over IPv6 MUST support [RFC 2890].

The sRouter MUST support the TR-181 GRE data model elements found in [TR-181].

An sRouter that supports GRE tunneling and SNMP MUST support the [CLAB-GRE-MIB].

When the sRouter is provisioned via the GRE tunneling MIB Annex A.4, it MUST permit the index of the Wi-Fi SSID to be set to the index numbers defined in Annex A.1.

When the sRouter is provisioned via the [TR-181] GRE tunneling profile, it MUST permit the index of the Wi-Fi SSID to be set to the index numbers defined in Annex A.2.

Annex A SNMP MIB Objects supported by the sRouter (Normative)

This Annex defines the SNMP MIBs that the sRouter is required to implement if SNMPv2c is supported.

The sRouter MUST support the following MIB objects:

- Access Network Independent Dev MIB [CLAB-ANI-DEV-MIB]
- The System Group [RFC 3418]
- The SNMP Group [RFC 3418]
- The ifGeneralInformationGroup, the linkUpDownNotificationsGroup, and one of (ifFixedLengthGroup, ifHCFixedLengthGroup, ifPacketGroup, ifHCPacketGroup, ifVHCPacketGroup) [RFC 2863]
- All mandatory groups of ipMIBCompliance2 in the IP-MIB [RFC 4293]
- All mandatory groups of ipForwardReadOnlyCompliance in the IP-Forward-MIB [RFC 4292]
- All mandatory groups of udpMIBCompliance2 in the UDP-MIB [RFC 4113]
- All mandatory groups of tcpMIBCompliance2 in the TCP-MIB [RFC 4022]

A.1 sRouter Interface Numbering

The sRouter MUST use in its MIB tables, when appropriate, an ifIndex number of '1' for the Operator-Facing Interface and an ifIndex number of '2' for the first Customer-Facing Interface. The sRouter MUST use an ifIndex number in accordance with Table 7 for any additional Customer-Facing Interfaces.

Table 7 - sRouter Interface Numbering

Interface	Type
1	Access Network interface (Router Operator-Facing Interface)
2 - 4	Reserved
5 - 15	Ethernet interfaces
16 - 31	Reserved
32 - 39	USB interfaces
40 – 47	MoCA interfaces
48 - 199	Reserved
200 – 299	Customer-Facing IP interfaces
300 – 399	Operator-Facing IP interfaces
400 – 499	GRE tunnel interfaces
1xxyy	Wi-Fi and SSID interfaces (where xx corresponds to the Wi-Fi radio interface (0 – 99), and yy corresponds to the SSID logical interface for Wi-Fi radio xx with yy in the range 1 – 99)
500 - 599	Additional Router CPE interfaces
600 - 699	Router internal interfaces (optional)

sRouter devices that include one or more Wi-Fi radios MUST follow the interface numbering and naming conventions specified in section 6.2.1 of [WI-FI MGMT].

A.2 sRouter ifTable Requirements

The sRouter MUST implement the row entry specified in Table 8 for the ifTable as specified in [RFC 2863].

Table 8 - sRouter ifTable Row Entries

ifTable ([RFC 2863])	Row Entry
IfIndex	1
ifDescr	**“**eRouter Operator-Facing Interface”
IfType	other(1)
IfMtu	0
IfSpeed	0
ifPhysAddress	eRouter MAC address
IfAdminStatus	up(1)
ifOperStatus	up(1)
IfLastChange	per [RFC 2863]
ifInOctets	0
IfInNUCastPkts	Deprecated
IfInDiscards	0
IfInErrors	0
IfUnknownProtos	0
ifOutOctets	0
ifOutUCastPkts	0
IfOutNUCastPkts	Deprecated
IfOutDiscards	0
IfOutErrors	0
IfOutQlen	Deprecated
IfSpecific	Deprecated

Additionally, for all interfaces supported the sRouter MUST use the values contained in Table 9 for the referenced objects in the ifTable. This includes modifications to the MAX-ACCESS for specific objects, which differs from [RFC 2863] in some cases.

Table 9 - sRouter ifTable Row Entries for Supported Interfaces

ifTable Object	Ethernet	USB	MoCA	CFI IP	OFI IP	GRE	Wi-Fi	SSID
IfIndex	5 - 15	30 - 39	40 - 47	200 - 299	300 - 399	400 - 499	Per Annex A.1	Per Annex A.1
ifDescr	-	-	-	-	-	-	-	-
IfType	ethernetCsmacd	usb	moca	ipForward	ipForward	tunnel	ieee80211(71)	ieee80211(71)
IfMtu	-	-	-	-	-	-	-	-
IfSpeed	-	-	-	-	-	-	-	-
ifPhysAddress	Ethernet physical address	USB physical address	MoCA physical address	CFI IP MAC address	OFI IP MAC address	MAC associated with tunnel endpoint	Empty string	SSID physical address
IfAdminStatus	Per [RFC 2863]	Per [RFC 2863]	Per [RFC 2863]	Per [RFC 2863] implemented as read-only	Per [RFC 2863] implemented as read-only	Per [RFC 2863] implemented as read-only	Per [RFC 2863]	Per [RFC 2863]
ifOperStatus	Per [RFC 2863]	Per [RFC 2863]	Per[RFC 2863]	Per [RFC 2863]	Per [RFC 2863]	Per [RFC 2863]	Per [RFC 2863]	Per [RFC 2863]
IfLastChange	unspecified	unspecified	unspecified	unspecified	unspecified	unspecified	unspecified	unspecified
ifInOctets	-	-	-	-	-	-	-	-
IfInNUCastPkts	deprecated	deprecated	deprecated	deprecated	deprecated	deprecated	deprecated	deprecated
IfInDiscards	-	-	-	-	-	-	-	-
IfInErrors	-	-	-	-	-	-	-	-
IfUnknownProtos	-	-	-	-	-	-	-	-
ifOutOctets	-	-	-	-	-	-	-	-
ifOutUCastPkts	-	-	-	-	-	-	-	-
IfOutNUCastPkts	deprecated	deprecated	deprecated	deprecated	deprecated	deprecated	deprecated	deprecated
IfOutDiscards	-	-	-	-	-	-	-	-
IfOutErrors	-	-	-	-	-	-	-	-
IfOutQlen	deprecated	deprecated	deprecated	deprecated	deprecated	deprecated	deprecated	deprecated
IfSpecific	deprecated	deprecated	deprecated	deprecated	deprecated	deprecated	deprecated	deprecated

A.3 sRouter ipNetToPhysicalTable Requirements

The sRouter MUST implement the row entry specified in Table 10 for the ipNetToPhysicalTable as specified in [RFC 4293] .

Table 10 - sRouter ipNetToPhysicalTable Row Entries

ipNetToPhysicalTable ([RFC 4293])	Router device
ipNetToPhysicalIfIndex	1
ipNetToPhysicalNetAddressType	ipv4(1) or ipv6(2)
ipNetToPhysicalNetAddress	Router IP Address
ipNetToPhysicalPhysAddress	Router MAC Address
ipNetToPhysicalLastUpdated	(refer to [RFC 4293])
ipNetToPhysicalType	static(4)
ipNetToPhysicalState	(refer to [RFC 4293])
ipNetToPhysicalRowStatus	'active'

A.4 CLAB-GRE-MIB

An sRouter that implements GRE tunneling MUST support the CableLabs GRE MIB [CLAB-GRE-MIB].

A.5 CLAB-MAP-MIB

The CableLabs Mapping of Address and Port MIB module was derived from the TR-181 data models for MAP-E and MAP-T. The full MIB text can be found at [CLAB-MAP-MIB].

A.6 CLAB-WIFI-MIB

If an sRouter supports an [802.11] interface and implements SNMP, the sRouter MUST support the CableLabs Wireless MIB [CLAB-WIFI-MIB].

A.7 CLAB-ANI-DEV-MIB

An sRouter that implements SNMP must support the [CLAB-ANI-DEV-MIB].

Annex B Provisioning and Operational Event Messages (Normative)

This list of the following locally logged Event Messages will facilitate resolution of issues. Access to the event log is vendor implementation specific.

Device	Error Code	Event ID	Severity	Intf	Event Message Text	Variables	Message Counter	Time Stamp
sRouter	H10.1	72001001	Informational		OFI - DHCPv4 Provisioning Complete	NA	NA	<time>
sRouter	H10.2	72001002	Informational		OFI - DHCPv6 Provisioning Complete	NA	NA	<time>
sRouter	H10.3	72001003	Critical		OFI - DHCPv4 Provisioning - X Retries attempted; Last attempt at <time>	NA	<count>	<time>
sRouter	H10.4	72001004	Critical		OFI - DHCPv6 Provisioning - X Retries attempted; Last attempt at <time>	NA	<count>	<time>
sRouter	H10.5	72001005	Critical		OFI - ICMPv6 No RA message received in response to RS	NA	<count>	<time>
sRouter	H10.6	72001006	Critical		OFI - ICMPv6 RA not properly configured for DHCPv6 (M = 0)	NA	NA	<time>
sRouter	H10.7	72001007	Critical		OFI - ICMPv6 Link Local DAD issue - Duplicate IP address detected	NA	NA	<time>
sRouter	H10.8	72001008	Critical		OFI - DHCPv4 No Offer / Ack message received	NA	<count>	<time>
sRouter	H10.9	72001009	Critical		OFI - DHCPv6 No Advertise / Reply message received	NA	<count>	<time>
sRouter	H10.10	72001010	Critical		OFI - DHCPv4 Missing Required DHCP option	<option #>	NA	<time>
sRouter	H10.11	72001011	Critical		OFI - DHCPv6 Missing Required DHCP option	<option #>	NA	<time>
sRouter	H10.12	72001012	Critical		OFI - DHCPv4 Bad value in required DHCP option	<option #>	NA	<time>
sRouter	H10.13	72001013	Critical		OFI - DHCPv6 Bad value in required DHCP option	<option #>	NA	<time>
sRouter	H10.14	72001014	Critical		OFI - DHCPv6 failed - No Address Available	NA	NA	<time>
sRouter	H10.15	72001015	Critical		OFI - DHCPv6 failed - No Prefix Available	NA	NA	<time>
sRouter	H10.16	72001016	Critical		OFI - DHCPv6 GUA DAD issue - Duplicate IP address detected	NA	NA	<time>
sRouter	H10.17	72001017	Critical		OFI - DHCPv6 Failure to renew Address	NA	<count>	<time>
sRouter	H10.18	72001018	Critical		OFI - DHCPv6 Failure to renew Prefix	NA	<count>	<time>
sRouter	H10.19	72001019	Critical		OFI - DHCPv4 Failure to renew lease	NA	<count>	<time>
sRouter	H10.20	72001020	Informational		OFI - DHCPv4 IP address released	NA	NA	<time>
sRouter	H10.21	72001021	Informational		OFI - DHCPv6 IP address released	NA	NA	<time>

Device	Error Code	Event ID	Severity	Intf	Event Message Text	Variables	Message Counter	Time Stamp
sRouter	H20.1	72002001	Critical		CFI - LAN Provisioning No Prefix available for sRouter interface(s)	<interface #>	NA	<time>
sRouter	H20.2	72002002	Critical		CFI - LAN Provisioning DHCPv6-PD No Prefix available - Subdelegation	NA	Client ID (duid)	<time>
sRouter	H20.3	72002003	Informational		CFI - LAN Provisioning DHCPv6-PD PD allocation mode set to width	NA	NA	<time>
sRouter	H20.4	72002004	Informational		CFI - LAN Provisioning DHCPv6-PD PD allocation mode set to depth	NA	NA	<time>
sRouter	H20.5	72002005	Informational		CFI - LAN Provisioning DHCPv6-PD PD hint of length 'X' received	X = [Null], Range [48 - 64]	Client ID (duid)	<time>
sRouter	H20.6	72002006	Alert		CFI - LAN Provisioning DHCPv6-PD Allocated prefix size X less than requested prefix Y - Subdelegation	X, Y	Client ID (duid)	<time>
sRouter	H20.7	72002007	Critical		CFI - LAN Provisioning DHCPv6 No client addresses available	NA	Client ID (duid)	<time>
sRouter	H20.8	72002008	Critical		CFI - Link Local DAD - Duplicate IP address detected	NA	Client ID (duid)	<time>
sRouter	H20.9	72002009	Critical		CFI - GUA DAD - Duplicate IP address detected	NA	Client ID (duid)	<time>
sRouter	H20.10	72002010	Critical		CFI - DHCPv6 Reconfigure Failure	NA	Client ID (duid)	<time>
sRouter	H20.11	72002011	Critical		CFI - LAN Provisioning DHCPv4 No Address available	NA	NA	<time>
sRouter	H20.12	72002012	Informational		CFI - LAN Provisioning DHCPv4 IP Address Conflict	NA	CHADDR	<time>
sRouter	H30.1	72003001	Informational		sRouter is administratively disabled	NA	NA	<time>
sRouter	H30.2	72003002	Informational		sRouter is enabled as IPv4 Only	NA	NA	<time>
sRouter	H30.3	72003003	Informational		sRouter is enabled as IPv6 Only	NA	NA	<time>
sRouter	H30.4	72003004	Informational		sRouter is enabled as Dual Stack	NA	NA	<time>

Annex C TR-069 Managed Objects Requirements (Normative)

The sRouter MUST support the objects associated with the Profiles and Components listed below. See [TR-106] for information about Components and Profiles in TR-069.

C.1 Profiles from [TR-181]

Table 11 - TR-181 Profiles for sRouter

Profile	Requirement	Notes
Download:1	MAY	
DownloadTCP:1	MAY	
Upload:1	MAY	
UploadTCP:1	MAY	
UDPEcho:1	MAY	
UDPEchoPlus:1	MAY	
SupportedDataModel:1	MUST	
SupportedDataModel:2	MUST	
MemoryStatus:1	MAY	
ProcessStatus:1	MAY	
TempStatus:1	MAY	
TempStatusAdv:1	MAY	
TempStatusAdv:2	MAY	
AutonXferComplPolicy:1	MAY	
User:1	MUST	
UPnPDev:1	MUST	Support Data model, other specs to detail UPNP functional requirements
UPnPDiscBasic:1	MUST	Support Data model, other specs to detail UPNP functional requirements
UPnPDiscAdv:1	MAY	
UPnPDiscAdv:2	MAY	
SelfTestDiag:1	MAY	
NSLookupDiag:1	MAY	
SimpleFirewall:1	MUST	
AdvancedFirewall:1	MUST	
USBHostsBasic:1	MAY	
USBHostsAdv:1	MAY	
PeriodicStatsBase:1	MAY	
PeriodicStatsAdv:1	MAY	
DownloadAnnounce:1	MAY	

Profile	Requirement	Notes
DownloadQuery:1	MAY	
Baseline:3	MUST	
Optical:1	MUST if interface supported	
EthernetRMONStats:1	MAY	
Ghn:1	MUST if interface supported	
DNSRelay:1	MAY	
Routing:1	MUST	
Routing:2	MUST	
IPv6Routing:1	MUST	
IPInterface:2	MUST	
IPv6Interface:1	MUST	
VLANTermination:1	MUST	
EthernetLink:1	MUST	
Bridge:1	MUST	
VLANBridge:1	MUST	
BridgeFilter:2	MUST	
BridgeFilterL3L4Filter:1	MAY	
EthernetInterface:1	MUST	
EthernetInterface:2	MAY	
ProviderBridge:1	MAY	
ProviderBridgeQoS:1	MAY	
LLDPBaseline:1	MAY	
LLDPRemOrgDeInfo:1	MAY	
GREBasic:1	MUST	
GREAdv:1	MUST	
MAPBasic:1	MUST	
MAPAdv:1	MUST	
HPNA:1	MUST if interface supported	
HPNADiagnostics:1	MUST if interface supported	
HPNAQoS:1	MUST if interface supported	
HomePlug:1	MUST if interface supported	
MoCA:1	MUST if interface supported	
UPA:1	MUST if interface supported	
UPDiagnostics:1	MUST if interface supported	

Profile	Requirement	Notes
WiFiRadio:1	Per [WI-FI MGMT]	
WiFiSSID:1	Per [WI-FI MGMT]	
WiFiAccessPoint:1	Per [WI-FI MGMT]	
WiFiEndPoint:1	Per [WI-FI MGMT]	
USBInterface:1	MUST if interface supported	
USBPort:1	MUST if interface supported	
NAT:1	MUST	
QoS:2	MAY	
QoSDynamicFlow:1	MAY	
QoSStats:1	MAY	
NeighborDiscovery:1	MUST	
RouterAdvertisement:1	MUST	
IPv6rd:1	MAY	
DSLite:1	MAY	
Hosts:2	MUST	
GatewayInfo:1	MUST	
DeviceAssociation:1	MUST	
UDPConnReq:1	MAY	
CaptivePortal:1	MAY	
Time:1	MAY	
IEEE8021xAuthentication:1	Per [WI-FI MGMT]	
IPPing:1	MAY	
TraceRoute:1	MAY	
DHCPv4Client:1	MUST	
DHCPv4Server:1	MUST	
DHCPv4CondServing:1	MAY	
DHCPv4Relay:1	MUST NOT	
DHCPv4ServerClientInfo:1	MUST	
DHCPv6Client:1	MUST	
DHCPv6ClientServerIdentity:1	MUST	
DHCPv6Server:1	MUST	
DHCPv6ServerAdv:1	MAY	
DHCPv6ServerClientInfo:1	MUST	
Processors:1	MAY	

Profile	Requirement	Notes
VendorLogFiles:1	MAY	
DUStateChngComplPolicy:1	MAY	
SM_ExecEnvs:1	MAY	
SM_DeployAndExecUnits:1	MAY	
SM_Baseline:1	MAY	
Location:1 Profile	MAY	
FaultMgmtSupportedAlarms:1 Profile	MAY	
FaultMgmtActive:1 Profile	MAY	
FaultMgmtHistory:1	MAY	
FaultMgmtExpedited:1 Profile	MAY	
FaultMgmtQueued:1	MAY	
BulkDataColl:1	MAY	
BullDataReports:1	MAY	
BulkDataJSONEncoding:1	MAY	
BulkDataCSVEncoding:1	MAY	
BulkDataHTTP:1	MAY	
BulkDataFileTransfser:1	MAY	
BulkDataStreaming:1	MAY	
IPsec:1	MAY	
IPsecAdv:1	MAY	
DNS_SD:1 Profile	MUST	
StandbyPolicy:1	MAY	
XMPPBasic:1 Profile	MAY	
XMPPConnReq:1	MAY	
XMPPAdvanced:1	MAY	
XMPPReconnect:1 Profile	MAY	
UDPEchoDiag:1	MAY	
ServerSelectionDiag:1	MAY	
InformParameters:1	MAY	
IEEE1905Device:1	MUST if interface supported	
IEEE1905TopologyMetric:1	MUST if interface supported	
IEEE1905TopologyNeighbor:1	MUST if interface supported	
IEEE1905TopologyHigherLayer	MUST if interface supported	
IEEE1905Power:1	MUST if interface supported	

Profile	Requirement	Notes
IEEE1905InterfaceSelection:1	MUST if interface supported	
IEEE1905LinkMetric:1	MUST if interface supported	
IEEE1905NetworkTopology:1	MUST if interface supported	
Device.PCP	MAY	MUST if DSLite is implemented

C.2 Extensions to TR-181 Profiles

The following are the CableLabs extensions to the profiles defined in [TR-181].

C.2.1 GRE Tunneling Extensions

The following are the CableLabs extensions to the profiles defined in [TR-181] for GRE tunneling.

Table 12 - CableLabs Extensions to TR-181 Profiles for GRE

Attribute Name	Type	Access	Type Constraints	Units	Default
KeepAliveCount	unsignedint	W			3
KeepAliveInterval	unsignedint	W		Seconds	60
KeepAliveFailureInterval	unsignedint	W		Seconds	300
KeepAliveRecoverInterval	unsignedint	W		Seconds	43200
MSSClampingValue	unsignedint	W	Disabled(0) Clamped(1) Clamping Size(>1)		0
ConcentratorServiceName	unsignedInt	W		FQDN	
RemoteEndpointConnectivityState	String(256)	RW			

An sRouter that implements GRE and [TR-069] MUST support the data objects in Table 12.

The GRE data object descriptions are as follows:

- KeepAliveCount – The number of keep-alive messages sent in a burst at regular intervals.
- KeepAliveInterval – Interval in seconds between keep-alive message bursts.
- KeepAliveFailureInterval – Time (in seconds) to wait after all available GRE concentrators fail to respond, before retrying the first GRE concentrator address.
- KeepAliveRecoverInterval – Time (in seconds) to remain on a secondary GRE concentrator, with clients connected, before retrying primary GRE concentrator. Zero value means no limit. Setting to a small non-zero value will cause an immediate switch from a secondary GRE concentrator back to the primary.
- MSSClampingValue – Specifies whether TCP MSS clamping is enabled on the tunnel. 0 disables clamping. 1 clamps the MSS depending on the interface MTU. A value > 1 will be used as clamping size.
- ConcentratorServiceName – FQDN of GRE tunnel concentrator/GW service. If this is set, then a DNS query of type SRV will be used for discovering the FQDN of remote endpoints on a GRE tunnel.
- RemoteEndpointConnectivityState – Comma-separated list (up to 4 items) of strings. Each item corresponds to one item in the RemoteEndpoints list, and contains one of the following strings: 'Reachable' indicates that the corresponding remote endpoint is responding to any configured KeepAlive messages. 'Unreachable' indicates that the remote endpoint has failed to adequately respond to the most recent KeepAlive attempt. 'NotInUse' indicates that the remote endpoint has not been used.

C.2.2 sRouter Extensions

The following are the CableLabs extensions to the Device.Routing.sRouter object defined in [TR-181].

Table 13 - CableLabs Extensions to TR-181 Profiles for sRouter

Attribute Name	Type	Access	Type Constraints	Units	Default
TopologyMode	boolean	R	depth(0) width(1)		
LinkIDEnable	boolean	R	disable(0) enabled(1)		Enabled

An sRouter MUST support the data objects in Table 13

C.3 Management Interface Protocols Requirements for GRE

Table 14 shows the mapping between the objects in the TR-181 data model and the SNMP MIB objects for GRE. CableLabs extension objects are included for completeness. An sRouter MUST support the data objects in Table 14.

Table 14 - GRE Data Model Objects

TR-181 Object Model	SNMP MIB Object	Requirement
Device.GRE		
TunnelNumberOfEntries	clabGRETunnelNumberOfEntries	Mandatory
FilterNumberOfEntries	clabGREFilterNumberOfEntries	Mandatory
Device.GRE.Tunnel.{i}.		
Enable	clabGRETunnelEnable	Mandatory
Status	clabGRETunnelStatus	Mandatory
Alias	clabGRETunnelAlias	Mandatory
RemoteEndpoints	clabGRETunnelRemoteEndpoints	Mandatory
KeepAlivePolicy	clabGRETunnelKeepAlivePolicy	Mandatory
KeepAliveTimeout	clabGRETunnelKeepAliveTimeout	Mandatory
KeepAliveThreshold	clabGRETunnelKeepAliveThreshold	Mandatory
DeliveryHeaderProtocol	clabGRETunnelDeliveryHeaderProtocol	Mandatory
DefaultDSCPMark	clabGRETunnelDefaultDscpMark	Mandatory
ConnectedRemoteEndpoint	clabGRETunnelConnectedRemoteEndpoint	Mandatory
InterfaceNumberOfEntries	clabGRETunnelInterfaceNumberOfEntries	Mandatory
X_CABLELABS_COM_KeepAliveCount	clabGRETunnelKeepAliveCount	Mandatory
X_CABLELABS_COM_KeepAliveInterval	clabGRETunnelKeepAliveInterval	Mandatory
X_CABLELABS_COM_KeepAliveFailureInterval	clabGRETunnelKeepAliveFailureInterval	Mandatory
X_CABLELABS_COM_KeepAliveRecoverInterval	clabGRETunnelKeepAliveRecoverInterval	Mandatory
X_CABLELABS_COM_MSSClampingValue	clabGRETunnelTcpMssClamping	Mandatory
X_CABLELABS_COM_ConcentratorServiceName	clabGRETunnelConcentratorServiceName	Mandatory
X_CABLELABS_COM_RemoteEndpointConnectivityState	clabGRETunnnelRemoteEndpointConnectivityState	Mandatory
Device.GRE.Tunnel.{i}.Stats.		
KeepAliveSent	clabGRETunnelStatsKeepAliveSent	Mandatory
KeepAliveReceived	clabGRETunnelStatsKeepAliveReceived	Mandatory
BytesSent	clabGRETunnelStatsBytesSent	Mandatory
BytesReceived	clabGRETunnelStatsBytesReceived	Mandatory
PacketsSent	clabGRETunnelStatsPacketsSent	Mandatory
PacketsReceived	clabGRETunnelStatsPacketsReceived	Mandatory
ErrorsSent	clabGRETunnelStatsErrorsSent	Mandatory

TR-181 Object Model	SNMP MIB Object	Requirement
ErrorsReceived	clabGRETunnelStatsErrorsReceived	Mandatory
Device.GRE.Tunnel.{i}.Interface.{i}.		
Enable	clabGRETunnelInterfaceEnable	Mandatory
Status	clabGRETunnelInterfaceStatus	Mandatory
Alias	clabGRETunnelInterfaceAlias	Mandatory
Name	clabGRETunnelInterfaceName	Mandatory
LastChange	clabGRETunnelInterfaceLastChange	Mandatory
LowerLayers	clabGRETunnelInterfaceLowerLayers	Mandatory
ProtocolIdOverride	clabGRETunnelInterfaceProtocolIdOverride	Mandatory
UseChecksum	clabGRETunnelInterfaceUseChecksum	Mandatory
KeyIdentifierGenerationPolicy	clabGRETunnelInterfaceKeyIdentifierGenerationPolicy	Mandatory
KeyIdentifier	clabGRETunnelInterfaceKeyIdentifier	Mandatory
UseSequenceNumber	clabGRETunnelInterfaceUseSequenceNumber	Mandatory
Device.GRE.Tunnel.{i}.Interface.{i}.Stats.		
BytesSent	clabGRETunnelInterfaceStatsBytesSent	Mandatory
BytesReceived	clabGRETunnelInterfaceStatsBytesReceived	Mandatory
PacketsSent	clabGRETunnelInterfaceStatsPacketsSent	Mandatory
PacketsReceived	clabGRETunnelInterfaceStatsPacketsReceived	Mandatory
ErrorsSent	clabGRETunnelInterfaceStatsErrorsSent	Mandatory
ErrorsReceived	clabGRETunnelInterfaceStatsErrorsReceived	Mandatory
DiscardChecksumReceived	clabGRETunnelInterfaceStatsDiscardChecksumReceived	Mandatory
DiscardSequenceNumberReceived	clabGRETunnelInterfaceStatsDiscardSequenceNumberReceived	Mandatory
Device.GRE.Filter.{i}.		
Enable	clabGREFilterEnable	Mandatory
Status	clabGREFilterStatus	Mandatory
Order	clabGREFilterOrder	Mandatory
Alias	clabGREFilterAlias	Mandatory
Interface	clabGREFilterInterface	Mandatory
AllInterfaces	clabGREFilterAllInterfaces	Mandatory
VLANIDCheck	clabGREFilterVlanIdCheck	Mandatory
VLANIDExclude	clabGREFilterVlanIdExclude	Mandatory
DSCPMarkPolicy	clabGREFilterDscpMarkPolicy	Mandatory

Annex D [RFC 6092] Simple Security Recommendations (Normative)

This section categorizes the recommendations from [RFC 6092] into recommendations for sRouter devices. While the RFC provides a good foundation for the development of a stateful inspection packet filtering firewall, it is not without omission and not all of its recommendations conform with best practices for cable networks. Additionally, the cable industry has developed several security mechanisms that supersede those provided in the recommendations. Where conflicts or recommendations other than those supplied by the RFC occur, they are called out explicitly.

D.1 Summary of Simple Security Requirements

This section provides a quick reference to the [RFC 6092] recommendations required by sRouter.

Critical - see Table 15
REC-3, REC-4, REC-5, REC-7, REC-10, REC-12, REC-14, REC-16, REC-18, REC-19, REC-21, REC-22, REC-23, REC-24, REC-25, REC-31, REC-32, REC-35, REC-36, REC-37, MSO-REC.

Important - see Table 16
REC-1, REC-2, REC-6, REC-8, REC-9, REC-11, REC-17, REC-33, REC-47.

BCP - see Table 17
REC-15, REC-20, REC-26, REC-27, REC-28, REC-29, REC-30, REC-38, REC-40, REC-41, REC-42, REC-43, REC-44, REC-45, REC-46, REC-48.

Other - see Table 18
REC-13, REC-49, REC-50.

Conflict - see Table 19
REC-34, REC-39.

D.2 Critical Recommendations

The sRouter MUST support [RFC 6092] on the Operator Facing interface. The following [RFC 6092] recommendations are critical to network connectivity and are to be included in all sRouter devices. All requirements in this section should be deemed mandatory as noted. These recommendations are in compliance with security requirements for the sRouter as the highest priority for development and testing.

Table 15 - Critical Recommendations

REC #	RFC 6092 Recommendation Text	Comments
REC-3	Packets bearing source and/or destination addresses forbidden to appear in the outer headers of packets transmitted over the public Internet MUST NOT be forwarded. In particular, site-local addresses are deprecated by [RFC 3879], and [RFC 5156] explicitly forbids the use of addresses with IPv4-Mapped, IPv4-Compatible, Documentation and ORCHID prefixes.	This would be the equivalent of an IPv6 bogon / martians list. Due to the CPU / memory resources of the devices and the fact that once deployed, it won't likely be changed, this would not include unallocated IPv6 space as it might on the backbone.
REC-4	Packets bearing deprecated extension headers prior to their first upper-layer-protocol header SHOULD NOT be forwarded or transmitted on any interface. In particular, all packets with routing extension header type 0 [RFC 2460] preceding the first upper-layer-protocol header MUST NOT be forwarded. (See [RFC 5095] for additional background.)	
REC-5	Outbound packets MUST NOT be forwarded if the source address in their outer IPv6 header does not have a unicast prefix assigned for use by globally reachable nodes on the interior network.	uRPF like behavior

REC #	RFC 6092 Recommendation Text	Comments
REC-7	By DEFAULT, packets with unique local source and/or destination addresses [RFC 4193] SHOULD NOT be forwarded to or from the exterior network.	Unique local addresses (ULA) can be forwarded between LAN interfaces on a customer premises router but as defined, should not exit the WAN interface. It is expected that ISP network will not carry routes for ULA address blocks so traffic will be dropped anyway.
REC-10	IPv6 gateways MUST forward ICMPv6 "Destination Unreachable" and "Packet Too Big" messages containing IP headers that match generic upper-layer transport state records.	If not, an MTU size mismatch can prevent connectivity, causing IPv6 sessions to fail. Conversely, if there is no state table entry, drop the packets.
REC-12	Filter state records for generic upper-layer transport protocols MUST NOT be deleted or recycled until an idle timer not less than two minutes has expired without having forwarded a packet matching the state in some configurable amount of time. By DEFAULT, the idle timer for such state records is five minutes.	If the timers are less than 2-5 minutes, many VPN tunnels break because the keep alive timer is often set to 360 seconds.
REC-14	A state record for a UDP flow where both source and destination ports are outside the well-known port range (ports 0-1023) MUST NOT expire in less than two minutes of idle time. The value of the UDP state record idle timer MAY be configurable. The DEFAULT is five minutes.	See REC-12 except this applies to low ports instead of high ports.
REC-16	A state record for a UDP flow MUST be refreshed when a packet is forwarded from the interior to the exterior, and it MAY be refreshed when a packet is forwarded in the reverse direction.	
REC-18	If a gateway forwards a UDP flow, it MUST also forward ICMPv6 "Destination Unreachable" and "Packet Too Big" messages containing UDP headers that match the flow state record.	Avoiding breaking path MTU discovery.
REC-19	Receipt of any sort of ICMPv6 message MUST NOT terminate the state record for a UDP flow.	If not supported, this could be employed in a DOS/DDOS attack against a CPE device by causing UDP sessions to close simply by receiving unsolicited ICMP reply messages.
REC-21	In their DEFAULT operating mode, IPv6 gateways MUST NOT prohibit the forwarding of packets, to and from legitimate node addresses, with destination extension headers of type "Authentication Header (AH)" [RFC 4302] in their outer IP extension header chain.	This requirement applies only to IPv6 packets. IPv4 IPsec AH packets should continue to be blocked from the Internet to internal hosts by default.
REC-22	In their DEFAULT operating mode, IPv6 gateways MUST NOT prohibit the forwarding of packets, to and from legitimate node addresses, with an upper-layer protocol of type "Encapsulating Security Payload (ESP)" [RFC 4303] in their outer IP extension header chain.	This requirement applies only to IPv6 packets. Hosts sufficient to support IPv6 should support rejecting unrequested AH/ESP packets by any hosts within the LAN/WAN. IPv4 IPsec AH packets should continue to be blocked from the Internet to internal hosts by DEFAULT.

REC #	RFC 6092 Recommendation Text	Comments
REC-23	If a gateway forwards an ESP flow, it MUST also forward (in the reverse direction) ICMPv6 "Destination Unreachable" and "Packet Too Big" messages containing ESP headers that match the flow state record.	
REC-24	In their DEFAULT operating mode, IPv6 gateways MUST NOT prohibit the forwarding of any UDP packets, to and from legitimate node addresses, with a destination port of 500, i.e., the port reserved by IANA for the Internet Key Exchange (IKE) Protocol [RFC 5996].	Blocking will likely break common L3 VPN (IPsec) connectivity. The IPsec IKE service listening on UDP/500 will not respond if it does not have a corresponding IPsec policy configured. As a result, leaving UDP/500 open could expose hosts to attack but could not be used in a reflection attack. IPv4 UDP/500 packets should continue to be blocked from the Internet to internal hosts by DEFAULT.
REC-25	In all operating modes, IPv6 gateways SHOULD use filter state records for Encapsulating Security Payload (ESP) [RFC 4303] that are indexable by a 3-tuple comprising the interior node address, the exterior node address, and the ESP protocol identifier. In particular, the IPv4/NAT method of indexing state records also by security parameters index (SPI) SHOULD NOT be used. Likewise, any mechanism that depends on detection of Internet Key Exchange (IKE) [RFC 5996] initiations SHOULD NOT be used.	ESP protocol identifier interactions may preclude more than one tunnel per endpoint.
REC-31	All valid sequences of TCP packets (defined in [RFC 793] MUST be forwarded for outbound flows and explicitly permitted inbound flows. In particular, both the normal TCP 3-way handshake mode of operation and the simultaneous-open mode of operation MUST be supported.	
REC-32	The TCP window invariant MUST NOT be enforced on flows for which the filter did not detect whether the window-scale option (see [RFC 1323]) was sent in the 3-way handshake or simultaneous-open.	Fast start support. May be more difficult for vendors, but could be necessary for high-latency connections.
REC-35	If a gateway cannot determine whether the endpoints of a TCP flow are active, then it MAY abandon the state record if it has been idle for some time. In such cases, the value of the "established flow idle-timeout" MUST NOT be less than two hours four minutes, as discussed in [RFC 5382]. The value of the "transitory flow idle-timeout" MUST NOT be less than four minutes. The value of the idle-timeouts MAY be configurable by the network administrator.	
REC-36	If a gateway forwards a TCP flow, it MUST also forward ICMPv6 "Destination Unreachable" and "Packet Too Big" messages containing TCP headers that match the flow state record.	Path MTU discovery and accessibility necessary for connectivity.
REC-37	Receipt of any sort of ICMPv6 message MUST NOT terminate the state record for a TCP flow.	This will prevent DoS against router due to unsolicited ICMPv6 messages.

REC #	RFC 6092 Recommendation Text	Comments
MSO-REC	By default an IGW MUST deny any protocol received on the WAN (operator facing) interface not specifically allowed by configuration with the following exceptions: DHCP, ND, ICMP and established TCP and UDP flows.	This recommendation is not found in [RFC 6092] but support is required in sRouter devices.

D.3 Important Recommendations

Failure to implement these [RFC 6092] recommendations could expose subscribers to infosec attacks. All sRouter implementations should support the list below as security requirements. The requirements below should be developed and tested after all critical requirements (Section D.2) are satisfied.

Table 16 - Important Recommendations

REC #	RFC 6092 Recommendation Text	Comments
REC-1	Packets bearing in their outer IPv6 headers multicast source addresses MUST NOT be forwarded or transmitted on any interface.	
REC-2	Packets which bear in their outer IPv6 headers multicast destination addresses of equal or narrower scope (see IPv6 Scoped Address Architecture [RFC 4007]) than the configured scope boundary level of the gateway MUST NOT be forwarded in any direction. The DEFAULT scope boundary level SHOULD be organization-local scope, and it SHOULD be configurable by the network administrator.	
REC-6	Inbound packets MUST NOT be forwarded if the source address in their outer IPv6 header has a global unicast prefix assigned for use by globally reachable nodes on the interior network.	Anti-spoofing
REC-8	By DEFAULT, inbound DNS queries received on exterior interfaces MUST NOT be processed by any integrated DNS resolving server.	Prevents DNS reflection attacks. It will also prevent subscribers from hosting a DNS server behind a router by default.
REC-9	Inbound DHCPv6 discovery packets [RFC 3315] received on exterior interfaces MUST NOT be processed by any integrated DHCPv6 server or relay agent.	Prevent recon scans (work around for vast IPv6 address space).
REC-11	If application transparency is most important, then a stateful packet filter SHOULD have "endpoint independent filter" behavior for generic upper-layer transport protocols. If a more stringent filtering behavior is most important, then a filter SHOULD have "address dependent filtering" behavior. The filtering behavior MAY be an option configurable by the network administrator, and it MAY be independent of the filtering behavior for other protocols. Filtering behavior SHOULD be endpoint independent by DEFAULT in gateways intended for provisioning without service-provider management.	For example, this would support allowing all http but reduces ability to block access to specific http websites since the solution uses the same port on the external interface. Since most gateways are not managed, that blocking is unlikely unless the device is subscribed to a reputation like service.

REC #	RFC 6092 Recommendation Text	Comments
REC-17	If application transparency is most important, then a stateful packet filter SHOULD have "endpoint-independent filtering" behavior for UDP. If a more stringent filtering behavior is most important, then a filter SHOULD have "address-dependent filtering" behavior. The filtering behavior MAY be an option configurable by the network administrator, and it MAY be independent of the filtering behavior for TCP and other protocols. Filtering behavior SHOULD be endpoint independent by DEFAULT in gateways intended for provisioning without service-provider management.	Similar to the REC-11 requirement but specific to UDP.
REC-33	If application transparency is most important, then a stateful packet filter SHOULD have "endpoint-independent filtering" behavior for TCP. If a more stringent filtering behavior is most important, then a filter SHOULD have "address-dependent filtering" behavior. The filtering behavior MAY be an option configurable by the network administrator, and it MAY be independent of the filtering behavior for UDP and other protocols. Filtering behavior SHOULD be endpoint independent by DEFAULT in gateways intended for provisioning without service-provider management.	Similar to the REC-11 requirement but specific to TCP.
REC-47	Valid sequences of packets bearing Shim6 payload extension headers in their outer IP extension header chains MUST be forwarded for all outbound and explicitly permitted flows. The content of the Shim6 payload extension header MAY be ignored for the purpose of state tracking.	

D.4 BCP Recommendations

The following [RFC 6092] recommendations are security best practices but are not critical to network communication. They may be supported as security requirements by sRouter devices, but are not deemed mandatory. These requirements should only be developed and tested after all requirements listed as critical (Section D.2) and important (Section D.3) have been implemented.

Table 17 - BCP Recommendations

REC #	RFC 6092 Recommendation Text	Comments
REC-15	A state record for a UDP flow where one or both of the source and destination ports are in the well-known port range (ports 0-1023) MAY expire after a period of idle time shorter than two minutes to facilitate the operation of the IANA-registered service assigned to the port in question.	This supports SIP, SKINNY, FTP or other parsers typically found in firewalls. By watching the control traffic, they can close a session early.
REC-20	UDP-Lite flows [RFC 3828] SHOULD be handled in the same way as UDP flows, except that the upper-layer transport protocol identifier for UDP-Lite is not the same as UDP; therefore, UDP packets MUST NOT match UDP-Lite state records, and vice versa.	UDP-Lite is an uncommon protocol and further implications may exist.
REC-26	In their DEFAULT operating mode, IPv6 gateways MUST NOT prohibit the forwarding of packets, to and from legitimate node addresses, with destination extension headers of type "Host Identity Protocol (HIP)" [RFC 5201] in their outer IP extension header chain.	Not currently a significant protocol, category approaches experimental.

REC #	RFC 6092 Recommendation Text	Comments
REC-27	The state records for flows initiated by outbound packets that bear a Home Address destination option [RFC 3775] are distinguished by the addition of the home address of the flow as well as the interior care-of address. IPv6 gateways MUST NOT prohibit the forwarding of any inbound packets bearing type 2 routing headers, which otherwise match a flow state record, and where A) the address in the destination field of the IPv6 header matches the interior care-of address of the flow, and B) the Home Address field in the Type 2 Routing Header matches the home address of the flow.	This will be needed to support IPv6 mobility but its use case in home is not clear at this time.
REC-28	Valid sequences of Mobility Header [RFC 3775] packets MUST be forwarded for all outbound and explicitly permitted inbound Mobility Header flows.	
REC-29	If a gateway forwards a Mobility Header [RFC 3775] flow, then it MUST also forward, in both directions, the IPv4 and IPv6 packets that are encapsulated in IPv6 associated with the tunnel between the home agent and the correspondent node.	
REC-30	If a gateway forwards a Mobility Header [RFC 3775] flow, then it MUST also forward (in the reverse direction) ICMPv6 "Destination Unreachable" and "Packet Too Big" messages containing any headers that match the associated flow state records.	
REC-38	All valid sequences of SCTP packets (defined in [RFC 4960]) MUST be forwarded for outbound associations and explicitly permitted inbound associations. In particular, both the normal SCTP association establishment and the simultaneous- open mode of operation MUST be supported.	If not implemented in first phase, SCTP should be dropped until this feature is implemented. Any unknown / unimplemented protocol MUST be dropped.
REC-40	If a gateway cannot determine whether the endpoints of an SCTP association are active, then it MAY abandon the state record if it has been idle for some time. In such cases, the value of the "established association idle-timeout" MUST NOT be less than two hours four minutes. The value of the "transitory association idle-timeout" MUST NOT be less than four minutes. The value of the idle-timeouts MAY be configurable by the network administrator.	
REC-41	If a gateway forwards an SCTP association, it MUST also forward ICMPv6 "Destination Unreachable" and "Packet Too Big" messages containing SCTP headers that match the association state record.	
REC-42	Receipt of any sort of ICMPv6 message MUST NOT terminate the state record for an SCTP association.	
REC-43	All valid sequences of DCCP packets (defined in [RFC 4340]) MUST be forwarded for all flows to exterior servers, and for any flows to interior servers with explicitly permitted service codes.	

REC #	RFC 6092 Recommendation Text	Comments
REC-44	A gateway MAY abandon a DCCP state record if it has been idle for some time. In such cases, the value of the "open flow idle-timeout" MUST NOT be less than two hours four minutes. The value of the "transitory flow idle- timeout" MUST NOT be less than eight minutes. The value of the idle-timeouts MAY be configurable by the network administrator.	
REC-45	If an Internet gateway forwards a DCCP flow, it MUST also forward ICMPv6 "Destination Unreachable" and "Packet Too Big" messages containing DCCP headers that match the flowstate record.	
REC-46	Receipt of any sort of ICMPv6 message MUST NOT terminate the state record for a DCCP flow.	
REC-48	Internet gateways with IPv6 simple security capabilities SHOULD implement a protocol to permit applications to solicit inbound traffic without advance knowledge of the addresses of exterior nodes with which they expect to communicate.	UPnP like functionality, but the protocol to do this reliably and the need to do this may not exist.

D.5 Other RFC 6092 Recommendations

These [RFC 6092] recommendations are not explicitly requirements for sRouter devices at this time. However, MSO consensus was reached for the incorporation of these requirements into sRouter to supplement and extend what is present in [RFC 6092]. These requirements should only be implemented after all other requirements have been satisfied.

Table 18 - Other 6092 Recommendations

REC #	RFC 6092 Recommendation Text	Comments
REC-13	Residential IPv6 gateways SHOULD provide a convenient means to update their firmware securely, for the installation of security patches and other manufacturer-recommended changes.	This requirement applies more to home routers owned by subscribers.
REC-49	Internet gateways with IPv6 simple security capabilities MUST provide an easily selected configuration option that permits a "transparent mode" of operation that forwards all unsolicited flows regardless of forwarding direction, i.e., not to use the IPv6 simple security capabilities of the gateway. The transparent mode of operation MAY be the default configuration.	The ability to turn off the firewall will probably be a requested feature but use case is still unclear as to who should be able to do this and when. The firewall MUST be on by default.

REC #	RFC 6092 Recommendation Text	Comments
REC-50	By DEFAULT, subscriber-managed residential gateways MUST NOT offer management application services to the exterior network.	Common management application services that need to be controlled include http (tcp/80), https (tcp/443), ssh (tcp/22), telnet (tcp/23) and snmp (udp161/162). As a default setting, it is important these be disabled to prevent blocking. Unclear impact to our ability to manage CPE devices like the integrated home gateway router. All externally facing management application services support authentication and require changing of all default credentials. All externally facing management application services also support restricting access to trusted IP blocks via an ACL. ACL(s) block both IPv4 and IPv6 by default unless explicitly allowed.

D.6 RFC 6092 Recommendations In Conflict With MSO Needs

The remaining recommendations from [RFC 6092] have been found to conflict with existing or proposed MSO requirements and should not be included in sRouter devices without explicit MSO approved modifications to render them useful. Such requirements should be interpreted to be "MUST NOT" to avoid such conflicts with MSO security policies.

Table 19 - RFC 6092 Recommendations In Conflict With MSO Needs

REC #	RFC 6092 Recommendation Text	Comments
REC-34	By DEFAULT, a gateway MUST respond with an ICMPv6 "Destination Unreachable" error code 1 (Communication with destination administratively prohibited), to any unsolicited inbound SYN packet after waiting at least 6 seconds without first forwarding the associated outbound SYN or SYN/ACK from the interior peer.	Preference would be to silently drop unsolicited packets from external sources rather than generate ICMPv6 unreachable due to administratively prohibited packets.
REC-39	By DEFAULT, a gateway MUST respond with an ICMPv6 "Destination Unreachable" error code 1 (Communication with destination administratively prohibited) to any unsolicited inbound INIT packet after waiting at least 6 seconds without first forwarding the associated outbound INIT from the interior peer.	Similar to syn dropping / errors. Prefer to silent drop instead of sending ICMP for DDoS protection and bounce attack protection.

Annex E sRouter GRE Tunneling Archtecture (Normative)

Figure 14 depicts a GRE forwarding model using one tunnel interface. It is included here to depict the data objects and indexing for mapping the interface, bridge ports and SSID when packets traverse the tunnel. The diagram differentiates between traffic via a private and public SSID. The configuration applies whether the configuration mechanism is SNMP or TR-069. For the purposes of this use case, it is presumed that provisioning of private and public SSIDs on the sRouter has already been completed. Data objects required for provisioning of the sRouter for private and public SSIDs is discussed elsewhere in this specification. It is also presumed that SSIDs 1 and 2 are public and SSIDs 3 and 4 are private.

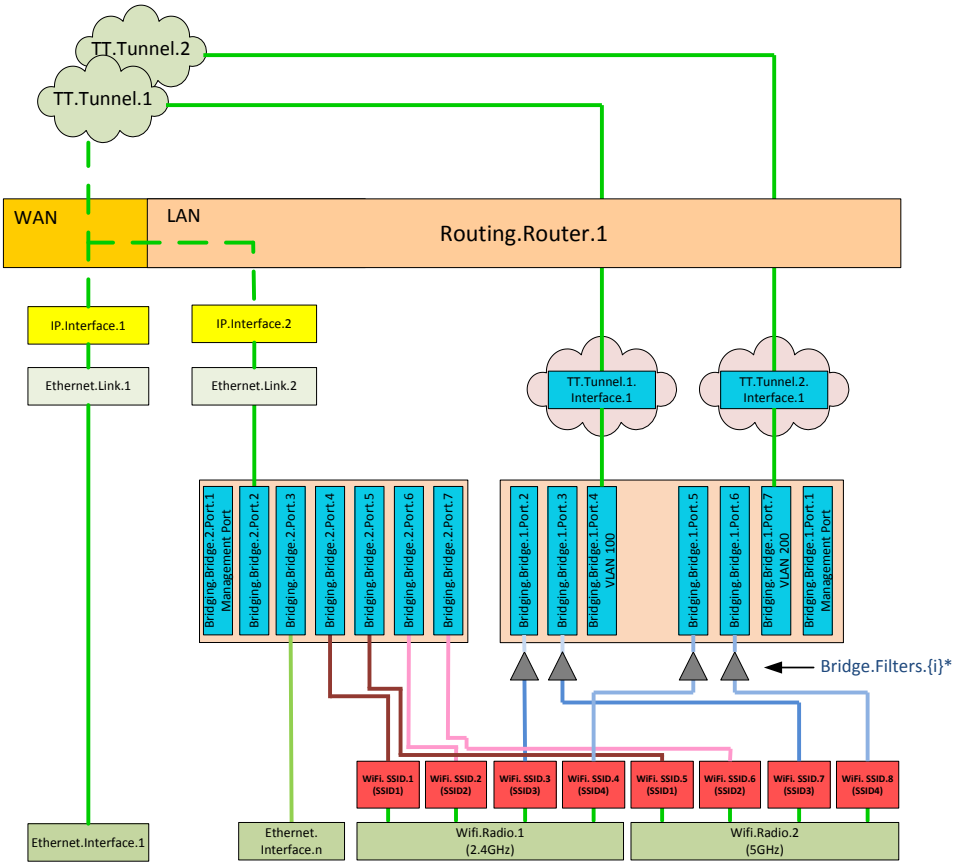

*Note: Classify and assign VLAN Tags for traffic separation.

Figure 14 - sRouter GRE Tunneling Architecture

Table 20 depicts the physical and logical interfaces, bridges and SSIDs that are referenced in Figure 14 above.

Table 20 - IF Indices and Row Instances for Data Objects Associated with GRE Tunneling

Row/Instance	Higher Layer Interface	if Index	Lower Layer	if Index
1	IP.Interface.1	300	Ethernet.Link.1	
2	Ethernet.Link.1		Ethernet.Interface.1	1
3	IP.Interface.2	200	Ethernet.Link.2	
4	Ethernet.Link.2		Bridging.Bridge.1.Port.4	

Row/Instance	Higher Layer Interface	if Index	Lower Layer	if Index
5	Bridging.Bridge.1.Port.1		Bridging.Bridge.1.Port.2 , Bridging.Bridge.1.Port.3, Bridging.Bridge.1.Port.4 , Bridging.Bridge.1.Port.5 , Bridging.Bridge.1.Port.6	
6	Bridging.Bridge.1.Port.2		Ethernet.Interface.2	
7	Bridging.Bridge.1.Port.3		WiFi.SSID.1 SSID1	10001
8	Bridging.Bridge.1.Port.4		WiFi. SSID.5 SSID1	10101
9	Bridging.Bridge.1.Port.5		WiFi. SSID.2 SSID2	10002
10	Bridging.Bridge.1.Port.6		WiFi.SSID.6 SSID2	10102
11	Bridging.Bridge.2.Port.1		Bridging.Bridge.2.Port.2 Bridging.Bridge.2.Port.3 Bridging.Bridge.2.Port.4 Bridging.Bridge.2.Port.5 Bridging.Bridge.2.Port.6 Bridging.Bridge.2.Port.7	
12	Bridging.Bridge.2.Port.2		WiFi.SSID.3 SSID3	10003
13	Bridging.Bridge.2.Port.3		WiFi.SSID.7 SSID3	10103
14	Bridging.Bridge.2.Port.5		WiFi.SSID.4 SSID4	10004
15	Bridging.Bridge.2.Port.6		WiFi.SSID.8 SSID4	10104
16	Bridging.Bridge.2.Port.4		TT.Tunnel.1.Interface.1	400
17	Bridging.Bridge.2.Port.7		TT.Tunnel.2.Interface.1	401
17	WiFi.SSID.1 WiFi.SSID.3 WiFi.SSID.5 WiFi.SSID.7	10001 10003 10101 10103	WiFi.Radio.1	10000
18	WiFi.SSID.2 WiFi.SSID.4 WiFi.SSID.6 WiFi.SSID.8	10002 10004 10102 10104	WiFi.Radio.2	10100
19	Wifi.Radio.1	10000		
20	Wifi.Radio.2	10100		

E.1 Use Case for Data Traffic Flow for Both Private and Public SSIDs

An sRouter that supports both private and public SSIDs must manage the private and public traffic separately. The following narrative describes how the private and public traffic traverses the physical and logical interfaces supported by the sRouter. There are four scenarios that will be described here: Private network outbound from the LAN, Private network inbound from the WAN, Public traffic outbound from a user on a public SSID, Public traffic inbound to a user on a public SSID. Each of these traffic flows is described below.

E.1.1 Private Network Outbound From the LAN

In this scenario, a user on the private network is connected via a private SSID and is attempting to connect to an outside network via the sRouter WAN interface. The following is an example of how the data traffic would flow.

E.1.2 Private Network Inbound From the WAN

In this scenario, traffic associated with a private SSID is routed through the sRouter WAN interface. The following is an example of how the data would flow from the Operator network through the internet gateway to the private Wi-Fi end-user device.

Traffic enters via Ethernet.Interface.1 and Ethernet.Link.1 to IP.Interface.1 (sRouter WAN interface). From here, it is routed to IP.Interface.2 (sRouter LAN interface) and to Ethernet.Link.2. Traffic is the directed to a logical bridge (Bridging.Bridge.2.Port.1) and bridged to the private SSID (WiFi.SSID.1) where it is passed to the private user via the WiFi.Radio.1 (2.4G).

Similarly, any other private LAN traffic (such as Ethernet, MoCA, etc.) would travel through the same logical bridge, sRouter, etc., as the private wireless traffic—the only difference is the Customer-Facing network interface type.

E.1.3 Community Wi-Fi User Outbound Via Public SSID

A public user connects via WiFi.Radio.2 (5G) and associates with WiFi.SSID.7. The traffic enters a logical bridge (Bridging.Bridge.1.Port.3) and is switched to another port within the bridge (Bridging.Bridge.1.Port.7). All ingress traffic is classified based on provider provisioned packet and/or port criteria. In this example, any ingress traffic via SSID3 will be tagged with an 802.1Q service tag VLAN ID 200 and bridged to the egress port appropriate egress port. The egress port is logically connected to the GRE tunnel interface (TT.Tunnel.1.Interface.1) that, in turn, is mapped to a tunnel instance (TT.Tunnel.1). The tunnel endpoint is the sRouter WAN interface (IP.Interface.1). Traffic exits the sRouter via Ethernet.Link.1 and Ethernet.Interface.1 (Operator-Facing interface).

NOTE: This example describes a single tunnel interface for a single user. However, there can be multiple users on a single tunnel interface, or multiple interfaces, depending upon a service operator's preference for managing traffic or vendor implementation. Such differentiation of traffic management is out of scope for this use case.

E.1.4 Community Wi-Fi User Inbound Via Public SSID

802.1Q service frames tagged with VLAN ID 200 and destined for the Community Wi-Fi end-user enters the gateway via the Operator-Facing interface and is passed to the Ethernet.Interface.1 and Ethernet.Link.1 to IP.Interface.1 (sRouter WAN interface). It is then placed on the tunnel (TT.Tunnel.1) and mapped to the logical tunnel interface (TT.Tunnel.1.Interface.1). Traffic then ingresses to the logical bridge (Bridging.Bridge.1.Port.7) and switched to (Bridging.Bridge.1.Port.3) where the VLAN tag is stripped and the frame is bridged to the public SSID (WiFi.SSID.7). Traffic is then transmitted to the user device using WiFi.Radio.2 (5G).

Annex F sRouter Certificate Public Key Infrastructure (Normative)

The sRouter security architecture uses an x509 digital certificate public key infrastructure (PKI) [RFC 5280] for authenticating endpoints on secure connections. This annex defines the PKI hierarchy, certificate profiles and CRL profiles that are used. CableLabs manages the PKI. All certification authority (CA) certificates are hosted by CableLabs or an authorized trusted 3rd party. Note: intermediate CA certificate Subject DN attributes may change due to maintenance/management of the PKI, which would also cause the Issuer DN attributes to change of end entity certificates. MSOs/operators should contact CableLabs to acquire management/provisioning server certificates (TR-069 ACS, SNMP, file server, GRE concentrator). sRouter vendors/manufacturers should contact CableLabs to acquire sRouter device certificates.

The profile tables below contain extension criticality settings, which are defined in [RFC 5280] as follows:

"A certificate-using system MUST reject the certificate if it encounters a critical extension it does not recognize or a critical extension that contains information that it cannot process. A non-critical extension MAY be ignored if it is not recognized, but MUST be processed if it is recognized."

This definition also applies to CRLs. This specification requires implementations to support critical extensions listed in the profile tables below. Support of non-critical extensions is optional unless stated otherwise in the specification.

The sRouter PKI hierarchy is shown below:

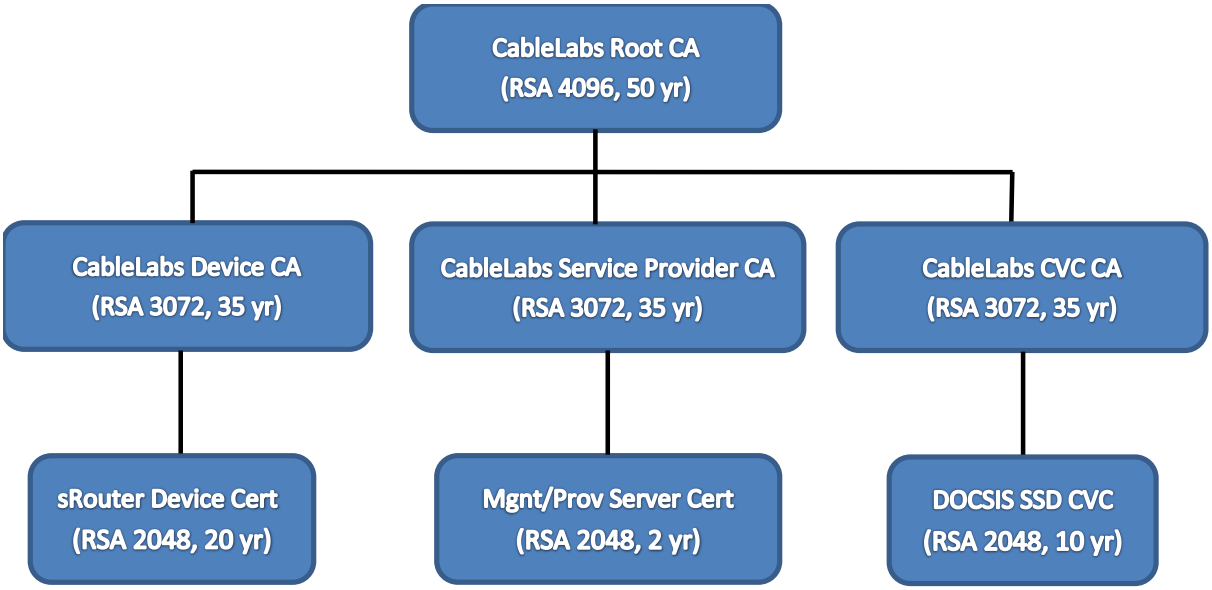

Figure 15 - Certificate Hierarchy

To help simplify and consolidate certificates, the root and intermediate CA certificates are shared with other projects (e.g., DOCSIS 3.1). End-entity certificates may also be shared with other CableLabs projects when serving the same function and having the same profile. For example, if a vendor has a certificate requesting account (CRA) for DOCSIS 3.1 CM device certificates, it may be possible to use the same CRA for issuing sRouter device certificates. Code verification certificates (CVCs) are used with DOCSIS SNMP secure software download. CVCs already issued and used for DOCSIS may also be used for signing sRouter images and download using the DOCSIS SNMP method.

All certificates and CRLs described in this annex are signed with the RSA signature algorithm, using SHA-256 as the hash function. The RSA signature algorithm is described in PKCS #1 [RSA 1]; SHA-256 is described in [FIPS 180-4].

F.1 CableLabs Root CA Certificate

Table 21 - CableLabs Root CA Certificate

Attribute Name		Settings		
Version		v3		
Serial number		Unique Positive Integer assigned by the CA		
Issuer DN		c=US o=CableLabs ou=Root CA01 cn=CableLabs Root Certification Authority		
Subject DN		c=US o=CableLabs ou=Root CA01 cn=CableLabs Root Certification Authority		
Validity Period		50 yrs		
Public Key Algorithm		Sha256WithRSAEncryption (1 2 840 113549 1 1 11)		
Keysize		4096 bits		
Parameters		NULL		
Standard Extensions	***OID***	***Include***	***Criticality***	***Value***
keyUsage	{id-ce 15}	X	TRUE	
keyCertSign				Set
cRLSign				Set
basicConstraints	{id-ce 19}	X	TRUE	
cA				Set
subjectKeyIdentifier	{id-ce 14}	X	FALSE	
keyIdentifier				Calculated per Method 1
subjectAltName	{id-ce 17}	O	FALSE	
directoryName				Set by the issuing CA

F.2 CableLabs Device CA Certificate

Table 22 - CableLabs Device CACertificate

Attribute Name		Settings		
Version		v3		
Serial number		Unique Positive Integer assigned by the CA		
Issuer DN		c=US o=CableLabs ou=Root CA01 cn=CableLabs Root Certification Authority		
Subject DN		c=US o=CableLabs ou=Device CA01 cn=CableLabs Device Certification Authority		
Validity Period		35 yrs		
Public Key Algorithm		Sha256WithRSAEncryption (1 2 840 113549 1 1 11)		
Keysize		3072-bits		
Parameters		NULL		
Standard Extensions	***OID***	***Include***	***Criticality***	***Value***
keyUsage	{id-ce 15}	X	TRUE	
keyCertSign				Set
cRLSign				Set
basicConstraints	{id-ce 19}	X	TRUE	
cA				Set
pathLenConstraint				0
subjectKeyIdentifier	{id-ce 14}	X	FALSE	
keyIdentifier				Calculated per Method 1
authorityKeyIdentifier	{id-ce 35}	X	FALSE	
keyIdentifier				Calculated per Method 1
subjectAltName	{id-ce 17}	O	FALSE	
directoryName				Set by the issuing CA for online CAs

F.3 sRouter Device Certificate

Table 23 - sRouter Device Certificate

Attribute Name		Settings		
Version		v3		
Serial number		Unique Positive Integer assigned by the CA		
Issuer DN		c=US o=CableLabs ou=Device CA01 cn=CableLabs Device Certification Authority		
Subject DN		c=<Country of Manufacturer> o=<Company Name> ou=<Manufacturing Location> cn= <MAC Address>		
Validity Period		20 yrs		
Public Key Algorithm		Sha256WithRSAEncryption (1 2 840 113549 1 1 11)		
Keysize		2048 bits		
Parameters		NULL		
Standard Extensions	***OID***	***Include***	***Criticality***	***Value***
keyUsage	{id-ce 15}	X	TRUE	
digitalSignature				Set
keyEncipherment				Set
authorityKeyIdentifier	{id-ce 35}	X	FALSE	
keyIdentifier				Calculated per Method 1

Values in angle brackets (<>) indicate that appropriate text as indicated below is present:

<Country of Manufacturer>: two-letter country code;

<Company Name>: name that identifies the company;

<Manufacturing Location>: name that identifies the location of manufacture;

<MAC Address>: MAC address of the sRouter.

The MAC Address is expressed as six pairs of hexadecimal digits separated by single colons (:), e.g., 00:60:21:A5:0A:23. Hexadecimal digits greater than 9 are expressed as uppercase letters.

F.4 CableLabs Service Provider CA Certificate

Table 24 - CableLabs Device CACertificate

Attribute Name	Settings
Version	v3
Serial number	Unique Positive Integer assigned by the CA
Issuer DN	c=US o=CableLabs ou=Root CA01 cn=CableLabs Root Certification Authority
Subject DN	c=US o=CableLabs ou=Service Provider CA01 cn=CableLabs Service Provider Certification Authority
Validity Period	35 yrs

Attribute Name		Settings		
Public Key Algorithm		Sha256WithRSAEncryption (1 2 840 113549 1 1 11)		
Keysize		3072 bits		
Parameters		NULL		
Standard Extensions	***OID***	***Include***	***Criticality***	***Value***
keyUsage	{id-ce 15}	X	TRUE	
keyCertSign				Set
cRLSign				Set
basicConstraints	{id-ce 19}	X	TRUE	
cA				Set
pathLenConstraint				0
subjectKeyIdentifier	{id-ce 14}	X	FALSE	
keyIdentifier				Calculated per Method 1
authorityKeyIdentifier	{id-ce 35}	X	FALSE	
keyIdentifier				Calculated per Method 1
subjectAltName	{id-ce 17}	O	FALSE	
directoryName				Set by the issuing CA for online CAs
authorityInfoAccess	{id-pe 1}	X	FALSE	Contains a single AccessDescription with an OCSP accessMethod and responder accessLocation HTTP URI.

F.5 Management/Provisioning Server Certificate (TR-069, SNMP, file server, GRE concentrator)

Table 25 - Management/Provisioning Server Certificate (TR-069, SNMP, file server, GRE concentrator)

Attribute Name		Settings		
Version		v3		
Serial number		Unique Positive Integer assigned by the CA		
Issuer DN		c=US o=CableLabs ou=Service Provider CA01 cn=CableLabs Service Provider Certification Authority		
Subject DN		c=<Country> o=<Company Name> cn=<server FQDN>		
Validity Period		5 yrs		
Public Key Algorithm		Sha256WithRSAEncryption (1 2 840 113549 1 1 11)		
Keysize		2048 bits		
Parameters		NULL		
Standard Extensions	***OID***	***Include***	***Criticality***	***Value***
keyUsage	{id-ce 15}	X	TRUE	
digitalSignature				Set
keyEncipherment				Set
authorityKeyIdentifier	{id-ce 35}	X	FALSE	
keyIdentifier				Calculated per Method 1
subjectAltName	{id-ce 17}	X	FALSE	
dNSName				<server FQDN>
authorityInfoAccess	{id-pe 1}	X	FALSE	Contains a single AccessDescription with an OCSP accessMethod and responder accessLocation HTTP URI.

F.6 CableLabs DOCSIS CVC CA Certificate

Table 26 - CableLabs DOCSIS CVC CA Certificate

Version	v3
Serial number	Unique Positive Integer assigned by the CA
Issuer DN	c=US o=CableLabs ou=Root CA01 cn=CableLabs Root Certification Authority
Subject DN	c=US o=CableLabs ou=CVC CA01 cn=CableLabs CVC Certification Authority
Validity Period	35 yrs
Public Key Algorithm	Sha256WithRSAEncryption (1 2 840 113549 1 1 11)
Keysize	3072-bits

Version		v3		
Parameters		NULL		
Standard Extensions	***OID***	***Include***	***Criticality***	***Value***
keyUsage	{id-ce 15}	X	TRUE	
keyCertSign				Set
cRLSign				Set
basicConstraints	{id-ce 19}	X	TRUE	
cA				Set
pathLenConstraint				0
subjectKeyIdentifier	{id-ce 14}	X	FALSE	
keyIdentifier				Calculated per Method 1
authorityKeyIdentifier	{id-ce 35}	X	FALSE	
keyIdentifier				Calculated per Method 1
subjectAltName	{id-ce 17}	O	FALSE	
directoryName				Set by the issuing CA for online CAs

F.7 Code Verification Certificate

Table 27 - Code Verification Certificate

Version		v3		
Serial number		Unique Positive Integer assigned by the CA		
Issuer DN		c=US o=CableLabs ou=CVC CA01 cn=CableLabs CVC Certification Authority		
Subject DN		c=<Country of Manufacturer> o=<Company Name> cn=Code Verification Certificate		
Validity Period		Up to 10 yrs		
Public Key Algorithm		Sha256WithRSAEncryption (1 2 840 113549 1 1 11)		
Keysize		2048-bits		
Parameters		NULL		
Standard Extensions	***OID***	***Include***	***Criticality***	***Value***
extKeyUsage	{id-ce 37}	X	TRUE	
codeSigning				Set
authorityKeyIdentifier	{id-ce 35}	X	FALSE	
keyIdentifier				Calculated per Method 1

Values in angle brackets (<>) indicate that appropriate text as indicated below is present:

<Country of Manufacturer>: two-letter country code;

<Company Name>: name that identifies the company

Co-signer CVCs will have a unique numeric value for the <Company Name>, which is assigned by CableLabs. The value is a printable string of eight hexadecimal digits. Each hexadecimal digit in the name is chosen from the ranges 0x30 to 0x39 or 0x41 to 0x46. The string 0x3030303030303030 is not assigned.

When a CVC is renewed, the orgname and validity start time of the old CVC will be used for the new CVC. The validity end time will be extended 10 years.

F.8 Certificate Revocation List (CRL)

Table 28- Certificate Revocation List

Attribute Name		Settings		
Signature Algorithm		Sha256WithRSAEncryption (1 2 840 113549 1 1 11)		
Issuer DN		<Subject DN of issuing CA>		
This Update Time		Time CRL was signed		
Next Update Time		1 month after time CRL was signed		
List of Revoked Certificates				
Extensions	***OID***	***Include***	***Criticality***	***Value***
issuingDistributionPoint	{id-ce 28}	X	TRUE	
distributionPoint				
generalName				HTTP URI to current CRL
authorityInformationAccess	{id-pe 1}	X	FALSE	
accessMethod				Id-ad-caIssuers
accessLocation				HTTP URI to signing CA certificate

Appendix I Example: Routing with Link ID (Informative)

This appendix provides example IP addressing and routing using Link ID as described throughout this document. The intention is to provide a reference example to aid in the proper application of Link ID for consistent multi-router packet forwarding without a routing protocol. In this example, an sRouter is provisioned with a /56 IPv6 prefix and (4) Customer-Facing IP Interfaces:

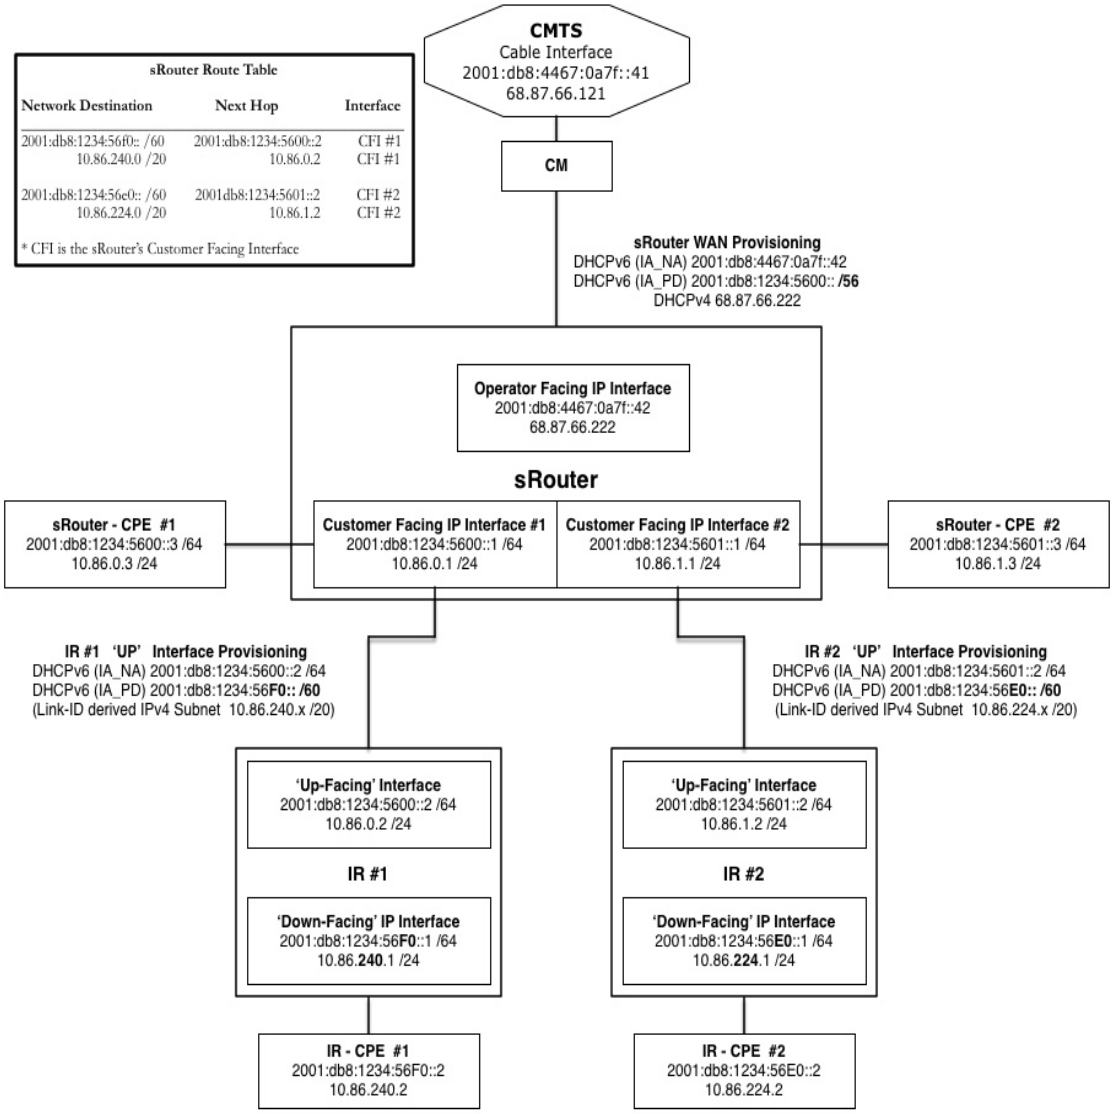

Figure 16 - Example of Link-ID with Prefix Delegation – Topology Mode Favors Width

```
sRouter Operator Facing Interface Provisioned Prefix: 2001:db8:1234:5600::/56

Customer-Facing IP Interface int(0)
     IPv6 Prefix:         2001:db8:1234:5600::/64
     Link ID Conversion:  56 00 -> 86 00
     IPv4 Network:   10.86.0.0/24   Gateway:  10.86.0.1
```

```
Customer-Facing IP Interface int(1)
     IPv6 Prefix:         2001:db8:1234:5600::/64
     Link ID Conversion:  56 00 -> 86 00
     IPv4 Network:   10.86.1.0/24 Gateway:  10.86.1.1

Customer-Facing IP Interface int(2)
     IPv6 Prefix:         2001:db8:1234:5601::/64
     Link ID Conversion:  56 01 -> 86 01
     IPv4 Network:   10.86.1.0/24 Gateway:  10.86.1.1

Customer-Facing IP Interface int(3)
     IPv6 Prefix:         2001:db8:1234:5601::/64
     Link ID Conversion:  56 01 -> 86 01
     IPv4 Network:   10.86.1.0/24 Gateway:  10.86.1.1

Topology Mode Encoding: Favor Width
Calculated on 3-Bit Boundary
Calculated subdelegation prefix length: /59
Number /59 available for prefix sub-delegation: 7
-- Sub-delegated prefix: 2001:db8:1234:5620::/59
-- Sub-delegated prefix: 2001:db8:1234:5640::/59
-- Sub-delegated prefix: 2001:db8:1234:5660::/59
-- Sub-delegated prefix: 2001:db8:1234:5680::/59
-- Sub-delegated prefix: 2001:db8:1234:56a0::/59
-- Sub-delegated prefix: 2001:db8:1234:56c0::/59
-- Sub-delegated prefix: 2001:db8:1234:56e0::/59

Sub-delegation:

The CER in this example has (4) Customer-Facing IP Interfaces - represented below as
int(0) -> int(3).
Assume there are (4) IRs with one attached to each of those Customer-Facing IP
Interfaces of the CER.
Further assume each IR is assigned an IPv6: IA_NA and IA_PD and IPv4 address in the
following way by the CER.

IR(0) on int(0) is  assigned via DHCPv6  IA_NA =  2001:db8:1234:9a00::100 and IA_PD =
2001:db8:1234:9a20::/59 and via DHCPv4 IP Address = 10.154.0.100
IR(1) on int(1) is  assigned via DHCPv6  IA_NA =  2001:db8:1234:9a01::100 and IA_PD =
2001:db8:1234:9a40::/59 and via DHCPv4 IP Address = 10.154.1.100
IR(2) on int(2) is  assigned via DHCPv6  IA_NA =  2001:db8:1234:9a02::100 and IA_PD =
2001:db8:1234:9a60::/59 and via DHCPv4 IP Address = 10.154.2.100
IR(3) on int(3) is  assigned via DHCPv6  IA_NA =  2001:db8:1234:9a03::100 and IA_PD =
2001:db8:1234:9a80::/59 and via DHCPv4 IP Address = 10.154.3.100

Route Table:

Network Destination               Next Hop                          Interface
-------------------------------------------------------------------------------
----------
2001:db8:1234:9a20::/59    2001:db8:1234:9a00::100    int(0)
10.154.32.0/19             10.154.0.100               int(0)
2001:db8:1234:9a40::/59    2001:db8:1234:9a01::100    int(1)
10.154.64.0/19             10.154.1.100               int(1)
2001:db8:1234:9a60::/59    2001:db8:1234:9a02::100    int(2)
10.154.96.0/19             10.154.2.100               int(2)
2001:db8:1234:9a80::/59    2001:db8:1234:9a03::100    int(3)
10.154.128.0/19            10.154.3.100               int(3)
```

I.1 IP MIB sRouter Example

The following table depicts a routing table example based on the Link ID reference example in Figure 16.

Route Description	DestType	Dest	PfxLen	Policy	NextHop Type	NextHop	IfIndex	RouteType	Metric (x)	RowStatus
LinkID IR#1 IPv4 CPE Network Route	ipv4 (1)	10.86.240.0	20	-	ipv4 (1)	10.86.0.2	CFI#1 IfIndex	remote(4)	-	active(5)
LinkID IR#1 IPv6 CPE Network Route	ipv6 (2)	2001:db8:1234:56f0::	60	-	ipv6 (2)	Link-Local of IR #1 'Up Facing' Interface	CFI#1 IfIndex	remote(4)	-	active(5)
LinkID IR#1 IPv4 CPE Network Route	ipv4 (1)	10.86.224.0	20	-	ipv4 (1)	10.86.1.2	CFI#2 IfIndex	remote(4)	-	active(5)
LinkID IR#2 IPv6 CPE Network Route	ipv6 (2)	2001:db8:1234:56e0::	60	-	ipv6 (2)	Link-Local of IR #2 'Up Facing' Interface	CFI#2 IfIndex	remote(4)	-	active(5)
Customer-Facing #1 IPv6 CPE Network Route	ipv6 (2)	2001:db8:1234:5600::	64	-	ipv6 (2)	Link-Local of CFI #1	CFI#1 IfIndex	remote(4)	-	active(5)
Customer-Facing #2 IPv6 CPE Network Route	ipv6 (2)	2001:db8:1234:5601::	64	-	ipv6 (2)	Link-Local of CFI #2	CFI#2 IfIndex	remote(4)	-	active(5)
Operating-Facing IPv4 Interface Host Route	ipv4 (1)	68.87.66.222	32	-	ipv4 (1)	-	OFI IfIndex	local (3)	-	active(5)
Customer-Facing #1 IPv4 Interface Host Route	ipv4 (1)	10.86.0.1	32	-	ipv4 (1)	-	CFI#1 IfIndex	local (3)	-	active(5)
Customer-Facing #2 IPv4 Interface Host Route	ipv4 (1)	10.86.1.1	32	-	ipv4 (1)	-	CFI#2 IfIndex	local (3)	-	active(5)

Route Description	DestType	Dest	PfxLen	Policy	NextHop Type	NextHop	IfIndex	RouteType	Metric (x)	RowStatus
Operating-Facing IPv6 Interface Host Route	ipv6 (2)	2001:db8:4467:0a7f::42	128	-	ipv6 (2)	-	OFI IfIndex	local (3)	-	active(5)
Customer-Facing #1 IPv6 Interface Host Route	ipv6 (2)	2001:db8:1234:5600::1	128	-	ipv6 (2)	-	CFI#1 IfIndex	local (3)	-	active(5)
Customer-Facing #2 IPv6 Interface Host Route	ipv6 (2)	2001:db8:1234:5601::1	128	-	ipv6 (2)	-	CFI#2 IfIndex	local (3)	-	active(5)
sRouter IPv4 Default Route	ipv4 (1)	0.0.0.0	0	-	ipv4 (1)	68.87.66.121	OFI IfIndex	remote (4)	-	active(5)
sRouter IPv6 Default Route	ipv6 (2)	::/0	0	-	ipv6 (2)	Link-Local of CMTS Cable Interface	OFI IfIndex	remote (4)	-	active(5)

Appendix II Acknowledgements (Informative)

On behalf of the cable industry and our member companies, CableLabs would like to thank the following individuals for their contributions to the development of this specification.

Contributor	Company Affiliation
Timothy Carey	Alcatel-Lucent
Kurt Lumbatis	Arris
Dan Torbet	Arris
Jason Schnitzer	Broadcom
Gordin Li	Broadcom
John McQueen	Broadcom
John Berg	CableLabs
Stuart Hoggan	CableLabs
Michael Kloberdans	CableLabs
Jon Schnoor	CableLabs
Volker Leisse	CableLabs
Matt Dillon	Charter
Erin Toney	Charter
Brent Bischoff	Cox
Baw Chng	Cox
Robert Schuler	Cox
Josh West	Cox
Ivan Lamoureux	Suddenlink
Kirk Erichsen	Time Warner
Pete Galt	Time Warner
Dale Hoeft	Time Warner
Daniel St. Pierre	Videotron
Peter Joyce	Virgin Media

Steve Burroughs, CableLabs

Appendix III Revision History

III.1 Engineering Changes for CL-SP-sRouter-I02-170111

The followings engineering changes are incorporated into CL-SP-sRouter-I02-170111:

ECN	Date Accepted	Summary	Author
sRouter-N-16.0147-2	5/19/2016	Add logging objects to CLAB-ANI-DEV-MIB	Luehrs, Kevin
sRouter-N-16.0155-2	12/1/2016	sRouter Security Requirements Update	Hoggan, Stuart